

SELECTIVE SIGN-DETERMINING MULTIPLE CONFIDENCE INTERVALS WITH FCR CONTROL

BY ASAF WEINSTEIN^{*,†} AND DANIEL YEKUTIELI[‡]

Stanford University[†] and Tel Aviv University[‡]

Given m unknown parameters with corresponding independent estimators, the Benjamini-Hochberg (BH) procedure can be used to classify the sign of parameters such that the expected proportion of erroneous directional decisions (directional FDR) is controlled at a preset level q . More ambitiously, our goal is to construct sign-determining confidence intervals – instead of only classifying the sign – such that the expected proportion of non-covering constructed intervals (FCR) is controlled. We suggest a valid procedure which employs the FCR-adjustment of Benjamini and Yekutieli [6] to a marginal confidence interval in order to construct a maximum number of sign-determining confidence intervals. The choice of the marginal confidence interval in our procedure controls a trade-off between power and the length of the constructed intervals. We propose a new marginal confidence interval that, when used in our procedure, allows to balance this trade-off and, in fact, often enjoy (almost) the best of both worlds.

We apply our methods to detect the sign of correlations in a highly publicized social neuroscience study and, in a second example, to detect the direction of association for SNPs with Type-2 Diabetes in GWAS data. In both examples we compare our procedure to existing methods and obtain encouraging results.

1. Introduction. Let f be a density symmetric about zero and suppose that an analyst collects independent observations $Y_i \sim f(y_i - \theta_i)$ $i = 1, \dots, m$ corresponding to unknown location parameters $\theta_i \in \mathbb{R}$. In many applications the analyst will highlight a subset of parameters which the data suggests as interesting, and then focus inference only on these highlighted parameters. A prototypical case is constructing confidence intervals for parameters θ_i corresponding to only rejected null hypotheses $H_{0i} : \theta_i = \theta_{0i}$, $i = 1, \dots, m$. For example, in RNA microarray one is interested in genes that are differentially expressed. More generally, we may consider a two-stage procedure, in which the analyst first attempts to answer a question of primary interest regarding each of the θ_i ; at the second stage, *only if* he was able to answer

^{*}Partially supported by the Simons foundation under the Math+X program

MSC 2010 subject classifications: Primary 62J15, 62F25

Keywords and phrases: Confidence intervals, Selective inference, Multiplicity, Directional decisions, False Coverage Rate, False Discovery Rate

the primary question regarding θ_i with enough certainty, the analyst will pose a secondary – follow-up – question regarding θ_i . The first question is often intended to detect “signals” of one or more types; the subsequent question may depend on the answer to the first, and is usually intended to learn about further qualities of θ_i .

In this article the primary question concerns the sign of the parameters. Specifically, the analyst is interested first in classifying the sign of parameters θ_i as positive (> 0) or non-positive (≤ 0); we refer to this as *weak* sign classification. A third decision – declaring “inconclusive data” – is allowed when an observation size is too small to infer the sign. The sign classification problem is important in many applications. For example, when comparing several drugs to a control it may be of interest to determine which are better (difference positive) and which are not (difference non-positive). Bohrer [7] and Bohrer and Schervish [8] called the sign classification problem a multiple *three-decision* problem (in reference to Neyman et al. [15]) and studied optimality of decision rules under family-wise error rate (FWER) control, namely, the probability of making at least one incorrect directional decision. In the current paper we consider rules that control the *weak directional* FDR, which we define as

$$(1.1) \quad \text{wdFDR} := \mathbb{E} \left[\frac{V_D}{R_D \vee 1} \right]$$

where R_D is the number of parameters whose sign was classified as positive or non-positive, and V_D is the number of parameters whose sign was incorrectly classified: non-positive parameters declared positive or positive parameter declared non-positive.

One procedure that is known to control wdFDR is the directional-BH procedure. Here parameters are selected via the usual BH procedure using two-sided p-values for testing $H_{0i} : \theta_i = 0$, and a directional decision is made for each selected parameter according to the sign of the estimator. The directional-BH procedure in fact makes a *strict* sign classification regarding each selected parameter, i.e., θ_i is declared negative (if $Y_i < 0$) or positive (if $Y_i > 0$), and still controls the expected proportion of incorrect decisions of any kind [6]. The latter is known as the *mixed directional* FDR [3], and is a stronger version of directional FDR; it may differ from the wdFDR when zero parameters are possible. The procedures we consider are required to make only weak directional decisions and control the wdFDR. To fix terminology, when we refer to a sign-classification or a directional decision, unless otherwise noted, it should be understood in the weak sense, i.e., positive or non-positive.

While a directional decision may be of primary importance, in practice it is almost always desirable to supplement such a decision with a confidence interval. For example, if the data suggests that a particular drug among a candidate set improves over control, we would like to be able to say how big the difference is at least; if the difference is immaterial, prescription of that drug may not be recommended after all. Suppose, then, that in addition to controlling wdFDR at the first stage, the analyst aims to construct a confidence interval for each parameter whose sign was classified. To match the error rate at the primary stage, we require the confidence intervals constructed at the follow-up stage to control the False Coverage Rate [6],

$$(1.2) \quad \text{FCR} = \mathbb{E} \left[\frac{V_{CI}}{R_{CI} \vee 1} \right]$$

where R_{CI} is the number of confidence intervals constructed and V_{CI} is the number of non-covering confidence intervals constructed. Since a confidence interval is constructed only for parameters that were reported as positive or as non-positive in the first stage, selection need be taken into account. Whereas without selection constructing marginal $1 - q$ level confidence intervals ensure $\text{FCR} \leq q$, Benjamini and Yekutieli [6] show that selection may dramatically increase not only the conditional non-coverage probability of each confidence interval, but even the average non-coverage rate over selected parameters, namely, the FCR.

One remaining issue is compatibility between the directional decision and the follow-up confidence interval. If a parameter θ_i was classified as non-positive (for example), it is undesirable that a supplementing confidence interval include positive values. Even if such a mishap did not lead to an actual contradiction (for example, if the error rates controlled at the two stages are not compatible), the analyst would find it inconvenient to report conflicting findings. Rosenblatt and Benjamini [19] raise this concern. Since controlling the wdFDR at the classification stage and the FCR at the confidence-interval stage does not automatically guarantee compatibility, we will add the requirement that a constructed confidence interval must agree with the directional decision: an interval including only positive values must be quoted if θ_i is classified as positive, and an interval including only non-positive values must be quoted if θ_i is classified as non-positive. A confidence interval that includes only positive or only non-positive values is said to *determine the sign* of the corresponding parameter.

We may now give a more formal definition of a two-stage procedure, which makes directional decisions and supplements them with compatible confidence intervals. To avoid cumbersome terminology, let us abbreviate

this by referring to simply a *two-stage* procedure. Formally, a two-stage procedure consists of

1. Selecting a subset $S = \mathcal{S}(Y) \subset \{1, \dots, m\}$
2. For each $i \in S$, classifying θ_i as positive or as non-positive
3. For each $i \in S$, constructing a confidence set C_i for θ_i such that $C_i \subseteq (0, \infty)$ if θ_i declared positive and $C_i \subseteq (-\infty, 0]$ if θ_i declared non-positive

We will say that a two-stage procedure is *valid at level q* if it has both $\text{wdFDR} \leq q$ and $\text{FCR} \leq q$. Of course, a valid two-stage procedure is not immediately desirable; for example, any sign-classification procedure which controls $\text{wdFDR} \leq q$ may be trivially extended to give a valid two-stage procedure by supplementing each classified parameter θ_i with the confidence interval $(0, \infty)$ or $(-\infty, 0]$ according as θ_i is classified as positive or non-positive. In this case $\text{FCR} = \text{wdFDR}$ and hence controlled at level q , but the confidence intervals provide no additional information. Designing valid procedures that furthermore tend to make many sign discoveries while constructing “informative” confidence intervals, will be of central interest in later sections; but the first question is how *valid* two-stage procedures which give nontrivial confidence intervals, can be constructed.

A natural strategy is to treat the two stages separately: at the first stage any classification rule controlling $\text{wdFDR} \leq q$ is specified; at the second stage confidence intervals are constructed according to *any* prescription that ensure $\text{FCR} \leq q$, treating the selection (classification) rule as given. Under independence of the Y_i , different approaches for achieving FCR control at the second stage are possible. Benjamini and Yekutieli [6] suggested to adjust the level of *marginal* confidence intervals so that FCR control is guaranteed. Alternatively, for “simple”¹ selection rules Zhong and Prentice [25] and Weinstein et al. [24] constructed confidence intervals that ensure *conditional* level- q coverage. Since the conditional approach restores nominal coverage for each interval given that it was selected, the FCR is indeed controlled [see, e.g., 24, Section 1]. Treating the two stages separately seems appealing, because the first stage can be determined by only considerations of maximizing the power of sign detection; while the second stage can take into account only considerations related to the shape (in particular, the length) of the confidence interval. However, when treating selection and inference separately we run the risk of constructing intervals that are incompatible

¹Informally, a “simple” selection rule means that for any $1 \leq i \leq m$, whenever i is selected the number of selected parameters is only a function of $\mathbf{Y}^{(i)} := (Y_1, \dots, Y_{i-1}, Y_{i+1}, \dots, Y_m)$.

with the directional decision.

In Section A.2 of the appendix we discuss connections to existing work on post-selection inference, and show that the conditional approach – that is, constructing a conditional $1 - q$ confidence interval for each selected parameter – indeed cannot ensure compatibility in general.

In the current article we suggest a procedure that reconciles directional decisions with follow-up confidence bounds. This is done by reversing the order in which the two stages are carried out: instead of deciding first on a selection (sign-classification) rule and adjusting follow-up intervals accordingly, we fix the confidence interval procedure and let it determine the set of parameters for which a directional decision is made. In the case of a single parameter, so that $Y \sim f(y - \theta)$, our procedure is very simple: suppose that $\mathcal{C}(y; \alpha)$ is any marginal $1 - \alpha$ confidence interval, i.e., $\Pr_{\theta}(\theta \notin \mathcal{C}(Y; \alpha)) \leq \alpha$. Then construct $\mathcal{C}(Y; \alpha)$ if and only if it is a subset of $(0, \infty)$ or of $(-\infty, 0]$. In that case,

$$(1.3) \quad \begin{aligned} \text{FCR} &= \Pr(\{\theta \notin \mathcal{C}(Y; \alpha)\} \cap \{\mathcal{C}(Y; \alpha) \text{ is constructed}\}) \\ &\leq \Pr(\theta \notin \mathcal{C}(Y; \alpha)) \leq \alpha. \end{aligned}$$

The sign of θ can be classified in the obvious way whenever a confidence interval is constructed, i.e., as positive or non-positive according as $\mathcal{C}(Y; \alpha)$ includes only positive or only non-positive values. Benjamini and Yekutieli [6] pointed out (1.3) to demonstrate that selection is handled gracefully as long as multiplicity is not involved. The main thrust of the current work is to extend the case of $m = 1$ to the case of general m . The idea is this: we would like to still allow the shape of the confidence interval at each Y_i determine whether it is constructed or not but, unlike in the single parameter case, selection need be taken into account. Indeed, the naïve procedure which, for some marginal confidence interval $\mathcal{C}(y; \alpha)$ constructs $\mathcal{C}(Y_i; q)$ whenever $\mathcal{C}(Y_i; q)$ includes only positive or only non-positive values, does not ensure $\text{FCR} \leq q$. Meanwhile, for any marginal confidence interval procedure, our method employs the FCR-adjustment of Benjamini and Yekutieli [6] to produce the largest set of *sign-determining* FCR-adjusted marginal confidence intervals. Since the choice of the *marginal* confidence interval procedure entirely determines our procedure, it controls both the power of the procedure as a sign-determining rule and the length (and shape) of constructed intervals. In line with recent work of Fithian et al. [12] and Tian and Taylor [22], we will see that in our procedure there is a trade-off between these objectives: higher power generally bears a cost of lower “accuracy” of the constructed intervals (as measured by their length and shape). This leads us to derive a new *marginal* confidence interval which enables our procedure to control the

trade-off. The corresponding procedure determines the sign of parameters according to a level- $(\psi \cdot 2q)$ directional-BH procedure for $1/2 \leq \psi \leq 1$, and constructs sign-determining intervals which, loosely speaking, are longer for larger ψ .

After studying the original sign problem, we offer two extensions of our procedure: the first is motivated by an example in genetics and generalizes our procedure to the case of a two-dimensional parameter, where the primary objective is to classify the sign of the first component. The second extension goes beyond the sign problem: we consider detecting parameters $\theta_i > \delta$ or $\theta_i < -\delta$ where δ is some pre-specified quantity. To address this problem we offer a procedure which constructs selective confidence intervals such that each interval contains either only values larger than δ or only values smaller than $-\delta$, while controlling the FCR.

The paper is organized as follows. Section 2 reviews the work of Benjamini and Yekutieli [6]. In Section 3 we give a formal definition of the type of procedures we are interested in, which we call Selective sign-determining CI procedures, and propose a class of valid procedures. Section 4 presents a new marginal confidence interval, designed specifically to be used in our selective sign-determining CI procedure. Results from a simulation study are reported in Section 5. In Section 6 we use our method to detect the sign of correlations in a neuroscience study. In Section 7 we extend our procedure to a two-dimensional case and apply the method to genomic data where we detect direction of association of SNPs with Type-2 diabetes. The appendix includes further simulation studies under dependency of the observations; an extension of the procedure for detecting only large correlations in the example of Section 6; and a discussion of related existing work on selective inference, where we contrast the conditional approach with ours. Proofs are also generally deferred to the appendix.

Notation. $\mathcal{C}(y; \alpha)$ is a confidence interval that covers the true value with probability at least $1 - \alpha$, and should be understood as a function of both y and α unless the context suggests otherwise. To emphasize the dependency on α , we sometimes write $\{\mathcal{C}(\cdot; \alpha) : 0 \leq \alpha \leq 1\}$ instead and call it a confidence interval *procedure*, but we use the two notations quite interchangeably. Throughout, f denotes a probability density and F is the corresponding distribution function. We denote by c_α the $1 - \alpha$ quantile of a distribution, that is, the value $F^{-1}(1 - \alpha)$; z_α is used for the special case of a standard Normal distribution. For convenience only, we write $Y_{(i)}$ for the observation with i -th *largest absolute value*. Finally, for any set B define $-B := \{-x : x \in B\}$. We tried to minimize the use of non-standard acronyms, but avoiding them altogether would result in a tedious manuscript. The few im-

portant ones are: CI=Confidence Interval; wdFDR=Weak Directional-FDR; SDCI=Sign-Determining Confidence Intervals; BY=Benjamini and Yekutieli [6]; QC=Quasi-Conventional; MQC=Modified Quasi-Conventional. The reader might find this list convenient to return to if any confusion arises.

2. Review. Benjamini and Yekutieli [6, BY hereafter] set up a framework for selective inference when multiple parameters are considered. Let $\mathbf{Y} = (Y_1, \dots, Y_m)$ be a vector of estimators where $Y_j \sim f(y_j - \theta_j)$. Suppose that \mathcal{S} is a pre-specified selection rule yielding a subset $S = \mathcal{S}(\mathbf{Y}) \subset \{1, \dots, m\}$ and that a procedure, which may depend on \mathcal{S} and on S , is used to construct confidence intervals for only the selected parameters $\{\theta_j : j \in S\}$. Denote by R_{CI} the number of confidence intervals constructed and by V_{CI} the number of non-covering confidence intervals constructed. Then BY define the false coverage-statement rate (FCR) to be the expected value of

$$Q_{CI} = \frac{V_{CI}}{R_{CI} \vee 1}$$

Thus the FCR depends on \mathcal{S} , which specifies what subset of parameters is selected in light of the data, and on the procedure which specifies how confidence intervals are constructed for any selected subset of parameters.

Suppose that at our disposal is a marginal confidence interval procedure $\{\mathcal{C}(\cdot; \alpha) : 0 \leq \alpha \leq 1\}$ which, for any $\alpha \in [0, 1]$, specifies a $(1 - \alpha)$ -level marginal confidence interval for θ based on $Y \sim f(y - \theta)$. That is, $\Pr_{\theta}(\theta \in \mathcal{C}(Y; \alpha)) \geq 1 - \alpha$ holds for any $\alpha \in [0, 1]$. Throughout the paper we will often refer to a marginal confidence interval *procedure* simply as a confidence *interval* and write $\mathcal{C}(y; \alpha)$, where it should be understood as a function of both y and α . Suppose that the procedure \mathcal{C} satisfies the following monotonicity requirement:

Requirement (MON 1) For any y and any $0 \leq \alpha, \alpha' \leq 1$, if $\alpha' \leq \alpha$ then $\mathcal{C}(y; \alpha) \subseteq \mathcal{C}(y; \alpha')$.

Denote $CI_i(\alpha) = \mathcal{C}(Y_i; \alpha)$. For $m > 1$, if \mathcal{S} is an arbitrary selection rule, then constructing $CI_i(q)$ for each $i \in S$ does not, in general, guarantee $\text{FCR} \leq q$. This should be obvious from considering, for example, a rule that selects the parameter corresponding to the largest of $m > 1$ independent estimators (here $R_{CI} \equiv 1$). On the other hand, constructing the marginal confidence interval at level $1 - q/m$ trivially ensures $\text{FCR} \leq q$. Indeed, denoting by NCI_i the event $\{\theta_i \notin CI_i(q/m), i \in S\}$, we have

$$\text{FCR} = \mathbb{E}[Q_{CI}] \leq \Pr(\cup_{i=1}^m NCI_i) \leq \Pr(\cup_{i=1}^m \{\theta_i \notin CI_i(q/m)\}) \leq q.$$

Yet under independence of the estimators, BY show that the Bonferroni adjustment is conservative, and a smaller increase in the confidence level is sufficient to ensure $\text{FCR} \leq q$. Specifically, they prove that the FCR is controlled at level q under the following scheme.

DEFINITION 1. Level- q BY FCR-Adjusted Selective-CI Procedure

1. Apply the selection criterion \mathcal{S} to obtain $\mathcal{S}(\mathbf{Y})$.
2. For each selected parameter θ_i , $i \in \mathcal{S}(\mathbf{Y})$, let

$$(2.1) \quad R_{\min}(\mathbf{Y}^{(i)}) = \min_y \left\{ \left| \mathcal{S}(\mathbf{Y}^{(i)}, Y_i = y) \right| : i \in \mathcal{S}(\mathbf{Y}^{(i)}, Y_i = y) \right\},$$

where $\mathbf{Y}^{(i)}$ is the vector obtained by omitting Y_i from \mathbf{Y} .

3. For each selected parameter θ_i , $i \in \mathcal{S}(\mathbf{Y})$, construct the following confidence interval:

$$CI_i \left(\frac{R_{\min}(\mathbf{Y}^{(i)}) \cdot q}{m} \right).$$

For many selection criteria, e.g., the step-up procedure of Benjamini and Hochberg, the term $|\mathcal{S}(\mathbf{Y}^{(i)}, Y_i = y)|$ is constant for all y such that $i \in \mathcal{S}(\mathbf{Y}^{(i)}, Y_i = y)$, implying that $R_{\min}(\mathbf{Y}^{(i)}) = R_{CI}$. In that case, to adjust the confidence intervals one simply multiplies the marginal non-coverage level q by the number of parameters selected and divides by m .

3. Selective-CI procedures that determine the sign. In this section we propose a general scheme to produce valid two-stage procedures, which relies on a marginal confidence interval: starting with any marginal confidence interval we show how a valid two-stage procedure can be obtained utilizing the FCR adjustment of Benjamini and Yekutieli [6]. We then turn to discuss how the choice of the marginal confidence interval affects the resulting selective procedure.

Recall the two-stage procedure described in the introduction. According to the definition, for a selected subset $S = \mathcal{S}(\mathbf{Y})$, the first stage consists of making a directional decision regarding each θ_i , $i \in S$; in the second stage a confidence interval which includes only positive or only non-positive values, is constructed for each θ_i , $i \in S$. It may have already occurred to the reader that the first stage is in fact redundant: there is no need for making directional decisions if the constructed intervals already determine the sign. To be precise, consider the following procedure instead.

DEFINITION 2. Selective Sign-Determining CI Procedure

1. Apply a pre-specified selection rule $\mathcal{S} : \mathcal{Y}^m \rightarrow 2^{\{1, \dots, m\}}$ to select a subset $S = \mathcal{S}(\mathbf{Y})$
2. For each $i \in \mathcal{S}(\mathbf{Y})$ construct a confidence interval CI_i for θ_i such that CI_i contains only positive or only non-positive values

A selective sign-determining CI (selective-SDCI hereafter) procedure is said to be valid at level q if it has $\text{FCR} \leq q$. Compared with a two-stage procedure, a selective-SDCI does not explicitly make directional decisions regarding selected parameters; however, we claim that nothing is compromised by neglecting that stage. Indeed, consider the extension of a selective-SDCI procedure to a two-stage procedure by classifying the sign of each $i \in S$ in the obvious way: if $CI_i \subseteq (0, \infty)$ declare θ_i as positive, and if $CI_i \subseteq (-\infty, 0]$ declare θ_i as non-positive. Then if an incorrect directional decision was made regarding θ_i , $i \in S$, it must be that $\theta_i \notin CI_i$. Therefore, for a selective-SDCI which is valid at level q ,

$$\begin{aligned} V_D &= \sum_{i=1}^m I(\theta_i \leq 0, \theta_i \text{ declared positive}) + I(\theta_i > 0, \theta_i \text{ declared non-positive}) \\ &\leq \sum_{i=1}^m I(\theta_i \notin CI_i) = V_{CI}, \end{aligned}$$

which implies that $\text{wdFDR} \leq \text{FCR} \leq q$. A Selective-SDCI is therefore just a more compact description of a two-stage procedure. Hence, from here on we will speak of a Selective-SDCI instead of a two-stage procedure.

We will now describe a general class of selective-SDCI procedures that control the FCR under independence of the estimators. Suppose that $\{\mathcal{C}(\cdot; \alpha) : 0 \leq \alpha \leq 1\}$ is any marginal confidence interval procedure satisfying Requirement (MON 1) of the previous section as well as

Requirement (MON 2) For any $0 \leq \alpha \leq 1$, $\mathcal{C}(-y; \alpha) = -\mathcal{C}(y; \alpha)$ and the lower boundary $l(y) = \inf \{\nu : \nu \in \mathcal{C}(y; \alpha)\}$ is increasing in $y > 0$.

Define a corresponding selective-SDCI procedure as follows.

DEFINITION 3. Level- q FCR-Adjusted Selective Sign-Determining CI Procedure

1. Let $Y_{(i)}$ be the estimate with the i -th *largest absolute value*, i.e., $|Y_{(m)}| \leq |Y_{(m-1)}| \leq \dots \leq |Y_{(1)}|$

2. Denoting $CI_i(\alpha) = \mathcal{C}(Y_i; \alpha)$, find

$$R = \max \left\{ r : CI_{(r)} \left(\frac{r \cdot q}{m} \right) \text{ does not include values of opposite signs} \right\}$$

and let $\mathcal{S}^*(\mathbf{Y}) = \{i : |Y_i| \geq Y_{(R)}\}$ be the (possibly empty) set of selected parameters.

3. For each $i \in \mathcal{S}^*(\mathbf{Y})$, construct the confidence interval

$$CI_i \left(\frac{R \cdot q}{m} \right).$$

THEOREM 1. *Suppose that $Y_i \sim f(y_i - \theta_i)$, $i = 1, \dots, m$ are independent and let $\{\mathcal{C}(\cdot; \alpha) : 0 \leq \alpha \leq 1\}$ be a marginal confidence interval procedure satisfying Requirement (MON 1) and Requirement (MON 2). Then the procedure in Definition 3 enjoys $\text{FCR} \leq q$.*

PROOF. We show that the FCR-adjusted selective-SDCI procedure uses the BY FCR-adjusted confidence level for the constructed CIs, in other words, the selective-SDCI procedure is just the BY procedure in Definition 1 for the selection rule \mathcal{S}^* in Definition 3. This will finish the proof, as the level- q BY procedure has $\text{FCR} \leq q$ for any selection rule.

It remains to show that for the procedure in Definition 3, $R_{\min}(\mathbf{Y}^{(i)}) = R$, in other words, $|\mathcal{S}^*(Y^{(i)}, Y_i = y)|$ is constant over y for all y such that $i \in \mathcal{S}^*(\mathbf{Y}^{(i)}, Y_i = y)$. Indeed, if this is true, then the constructed intervals use the BY-adjusted level and therefore $\text{FCR} \leq q$. This part is proved in the appendix. \square

For a given marginal confidence interval $\mathcal{C}(y; \alpha)$ the FCR-adjusted selective-SDCI procedure of Definition 3 constructs the largest number possible of BY FCR-adjusted confidence intervals that determine the sign. Since the set of discoveries is determined based on the adjusted marginal confidence intervals, our procedure is completely characterized by the choice of the marginal interval $\mathcal{C}(y; \alpha)$. Hence this choice affects both the tendency of the procedure to construct confidence intervals and the shape of the constructed confidence intervals: in particular, using a marginal interval $\mathcal{C}(y; \alpha)$ which determines the sign for relatively small values of $|y|$ will enhance selection, i.e., the construction of confidence intervals. Since, as discussed above, each constructed confidence interval may be associated with a directional decision, we informally refer to the tendency to construct many confidence intervals as the *power* of the selective-SDCI procedure. On the other hand, if a marginal interval $\mathcal{C}(y; \alpha)$ with relatively (to other marginal CIs) small maximum length

is used, then the constructed confidence intervals will also enjoy a relatively (to other marginal CIs adjusted at the same level) small maximum length, because only the confidence level is adjusted when constructing the intervals. The following examples describe the FCR-adjusted selective-SDCI procedure corresponding to three different choices of a marginal confidence interval. We will assume here that $Y_i - \theta_i$ are i.i.d. $N(0, 1)$ and denote $z_p = \Phi^{-1}(1 - p)$. As before, $Y_{(i)}$ is the estimate with the i -th largest absolute value.

- (a) *Symmetric confidence interval*. Set $\mathcal{C}(y; \alpha) = (y - z_{\alpha/2}, y + z_{\alpha/2})$. Since for any $\alpha \in (0, 1)$ this confidence interval includes values of one sign only (and possibly zero) whenever $z_{\alpha/2} \leq |y|$, the algorithm in Definition 3 selects the parameters corresponding to the $R = \max \{r : \{z_{r \cdot q/(2m)} \leq |Y_{(r)}|\}\}$ largest observations. Now let $P_i = 2(1 - \Phi(|Y_i|))$ be the two-sided p-value for testing $H_{0i} : \theta_i = 0$, and let $P_{(1)} \leq P_{(2)} \leq \dots \leq P_{(m)}$ be the ordered p-values (note that the subscript of the order statistic has the conventional meaning for the p-values but not for the estimators). Then $R = \max \{r : P_{(r)} \leq r \cdot q/m\}$ and so the selected parameters are exactly those corresponding to hypotheses rejected by the BH procedure applied at level q . The constructed confidence interval for each selected parameter θ_i is $CI_i = Y_i \pm z_{R \cdot q/(2m)}$.
- (b) *One-sided confidence interval*². Take

$$(3.1) \quad \mathcal{C}(y; \alpha) = \begin{cases} (-\infty, \infty), & -z_\alpha < y < z_\alpha \\ (0, \infty), & z_\alpha \leq y \\ (-\infty, 0], & y \leq -z_\alpha. \end{cases}$$

For any α this confidence interval includes values of one sign only already when $z_\alpha \leq |t|$. Our procedure therefore selects the set of parameters corresponding to the $R = \max \{r : z_{r \cdot q/m} \leq |Y_{(r)}|\} = \max \{r : P_{(r)} \leq r \cdot (2q)/m\}$ largest observations, which is the set of parameters rejected by the BH procedure when applied at level $2q$. The constructed confidence interval for each selected parameter θ_i is $CI_i = (0, \infty)$ if $0 < Y_i$ and $CI_i = (-\infty, 0]$ if $Y_i < 0$.

- (c) *Pratt's confidence interval*³. We can use a more sophisticated one-sided

²For lack of a better term we refer to the CI in (3.1) as “one-sided”, although this name is usually reserved for a CI of the form $(y - z_\alpha, \infty)$ or $(-\infty, y + z_\alpha)$

³The original CI suggested by Pratt treats zero “symmetricly”, whereas we append zero to the negative part of the line; (3.2) is therefore slightly different from the original construction, but the difference is not essential.

confidence interval,

$$(3.2) \quad \mathcal{C}(y; \alpha) = \begin{cases} (y - z_\alpha, y + z_\alpha), & \text{if } |y| < z_\alpha \\ (0, y + z_\alpha), & \text{if } z_\alpha \leq y \\ (y - z_\alpha, 0], & \text{if } y \leq -z_\alpha \end{cases}.$$

This construction was suggested by Pratt [18], who sought to minimize the expected length of a confidence interval at $\theta = 0$. Pratt's interval still determines the sign at z_α but its length is finite when it determines the sign, as opposed to the usual one-sided interval. The resulting FCR-adjusted selective-SDCI procedure therefore still has $R = \max \{r : z_{r,q/m} \leq |Y_{(r)}|\}$ and selects according to a level- $2q$ BH procedure. However, the constructed confidence interval for a selected parameter is now

$$CI_i = \begin{cases} (0, Y_i + z_{Rq/m}), & \text{if } z_{Rq/m} < Y_i \\ (Y_i - z_{Rq/m}, 0], & \text{if } z_{Rq/m} < Y_i \end{cases}$$

instead of the infinitely long intervals that are constructed with the plain one-sided interval (3.1).

It is easy to verify that all marginal confidence intervals above are valid (i.e., have $1 - \alpha$ coverage) and satisfy the two monotonicity requirements (MON 1) and (MON 2).

For a fixed α we would ideally want to equip the FCR-adjusted selective-SDCI procedure with a marginal interval $\mathcal{C}(y; \alpha)$ which determines the sign as early as possible, and at the same time has the smallest possible (say, maximum) length. Unfortunately, these two requests are incompatible: early sign determination has a price of longer confidence intervals, at least for some values of y . This is demonstrated in the examples above: the two-sided marginal interval has shortest possible maximum length, but determines the sign starting only at the $1 - \alpha/2$ quantile; whereas the one-sided marginal interval determines the sign already at the $1 - \alpha$ quantile, but has infinite maximum length. The Pratt interval improves on the length of the one-sided interval “for free”, but its length is still unbounded in y , which is necessary if sign determination starting at the $1 - \alpha$ quantile is desired. Consequently, if we are to use the procedure of Definition 3, then a trade-off between power and maximum (potential) length of the constructed intervals is unavoidable.

Nevertheless, we are not limited to the marginal confidence intervals in (a)-(c), in which sign determination occurs at either of the two extremes, $F^{-1}(1 - \alpha/2)$ or $F^{-1}(1 - \alpha)$. Instead of insisting on earliest possible sign determination or smallest possible maximum length, we may choose a marginal

confidence interval that balances between early sign determination and maximum length. That is, a marginal confidence interval which determines the sign starting at a value slightly bigger than the $1 - \alpha$ quantile, and in turn have maximum length that is only slightly larger than twice the $1 - \alpha/2$ quantile. Equipped with such a marginal family, the FCR-adjusted selective-SDCI procedure of Definition 3 will select parameters according to a BH procedure at a level close to $2q$, while controlling the length of the constructed confidence intervals.

Benjamini et al. [4] suggested a non-equivariant marginal confidence interval which is fit for the job. They assume that $Y \sim f(y - \theta)$ with $f = F'$ a unimodal, symmetric density, and obtain their Quasi-Conventional (QC hereafter) by inverting a family of acceptance regions. Specifically, the QC interval at y is defined as the convex hull of

$$\{\theta : y \in A_{QC}(\theta)\}$$

where

$$(3.3) \quad A_{QC}(\theta) = \begin{cases} (\theta - \bar{c}, \theta + \tilde{c}), & 0 < \theta \leq \bar{c} \\ (0, \theta + F^{-1}(1 - \alpha + F(-\theta))), & \bar{c} < \theta \leq c_{\alpha/2} \\ (\theta - c_{\alpha/2}, \theta + c_{\alpha/2}), & c_{\alpha/2} < \theta \end{cases}$$

and $A(\theta) = -A(-\theta)$ for $\theta < 0$. The acceptance region at zero is symmetric in the original construction but we take

$$A_{QC}(0) = (-\infty, c_\alpha)$$

which fits our (asymmetric) definition of sign determination. The constants \bar{c}, \tilde{c} are determined by a parameter $1/2 \leq \psi < 1$ and given by

$$\bar{c} = F^{-1}(1 - \psi\alpha) \quad \tilde{c} = F^{-1}(1 - \alpha + F(-\bar{c})).$$

For any $p \in [0, 1]$ we write $c_p = F^{-1}(1 - p)$ for the $(1 - p)$ -th quantile of F . The QC confidence interval determines the sign for $|y| \geq \bar{c} \in (c_\alpha, c_{\alpha/2}]$ and can be shown to have maximum length $\tilde{c} + c_{\alpha/2} < \infty$. The parameter ψ controls the balance between early sign determination and maximum length of the QC interval. For $\psi = 1/2$ we have $\bar{c} = c_{\alpha/2}$ and the usual symmetric confidence interval obtains. When $\psi \rightarrow 1$, $\bar{c} \rightarrow c_\alpha$ and for any fixed y the Pratt interval obtains in the limit. As ψ increases from $1/2$ to 1 , sign determination occurs at a gradually earlier point at the cost of an increasing maximum length.

Since for any α the QC interval with $1/2 \leq \psi < 1$ determines the sign at $\bar{c} < c_{\alpha/2}$, the FCR-adjusted selective-SDCI procedure using the QC interval

will have more power than using the symmetric interval. At the same time constructed intervals will be shorter as compared to using the Pratt confidence interval, and their length never exceeds $F^{-1}(1 - q' + F(-F^{-1}(1 - \psi q')))$ for $q' = Rq/m$ (this is just $\tilde{c} + c_{\alpha/2}$ for $q = q'$). While the QC interval already has the features that would make our procedure balance between power and length, an improvement is in fact possible. Indeed, we will show that the QC interval can be slightly modified so that our procedure constructs shorter intervals at no expense.

4. A Modified Quasi-Conventional CI. In this section we present a new marginal confidence interval that adopts a feature from Finner [11] to modify the QC interval of Benjamini et al. [4]. The idea is to take advantage of the fact that in the FCR-adjusted selective-SDCI procedure only sign-determining confidence intervals are ultimately constructed; hence inflating the QC confidence interval whenever it anyway includes values of opposite signs, has no cost on the one hand, and on the other hand it allows to construct shorter confidence intervals when the sign is determined.

As in Benjamini et al. [4] we make the further assumption that $f = F'$ is a unimodal density. We obtain the modified Quasi-Conventional (MQC hereafter) interval by modifying the acceptance regions (3.3). Hence, consider

$$(4.1) \quad A_{MQC}(\theta) = \begin{cases} (-\bar{c}, g(\theta)), & 0 < \theta \leq \bar{c} + c_{\alpha/2} \\ (\theta - c_{\alpha/2}, \theta + c_{\alpha/2}), & \bar{c} + c_{\alpha/2} < \theta \end{cases}$$

with $A_{MQC}(\theta) = -A_{MQC}(-\theta)$ for $\theta < 0$, and $A_{MQC}(0) = (-\infty, c_{\alpha})$. For $1/2 \leq \psi < 1$,

$$\bar{c} = F^{-1}(1 - \psi\alpha) \quad \tilde{c} = F^{-1}(1 - \alpha + F(-\bar{c})),$$

and

$$g(\theta) = \theta + F^{-1}\{1 - \alpha + F(-\bar{c} - \theta)\}.$$

As before, ψ is a parameter which controls how early the confidence interval determines the sign of θ , and is chosen in advance.

The MQC interval is obtained as the convex hull of $\{\theta : y \in A_{MQC}(\theta)\}$. Inverting the family of acceptance regions in (4.1) is more complicated than it is for the QC acceptance regions because we need to distinguish between three cases (i) $0 < \psi \leq \psi_1$ (ii) $\psi_1 < \psi \leq \psi_2$ and (iii) $\psi_2 < \psi$. Here

ψ_1 is the value of ψ such that $\tilde{c} = 2\bar{c} + c_{\alpha/2}$

ψ_2 is the value of ψ such that $\tilde{c} = \bar{c} + 2c_{\alpha/2}$.

From a practical point of view, however, at least when f is the standard Normal density, ψ_1 tends to be very close to 1. For example, when f is the standard Normal density and $\alpha = 0.1$, $\psi_1 > 0.999$, and it is even closer to 1 for smaller α . This means that for typical, small values of α , unless ψ is chosen extremely close to 1, we are in case (i) above. For clarity we present here the confidence bounds for the first case only, and defer the other two cases to the appendix. Hence, for $0 < \psi \leq \psi_1$, the convex hull of $\{\theta : y \in A_{MQC}(\theta)\}$ is given by

$$(4.2) \quad \mathcal{C}_{MQC}(y; \alpha) = \begin{cases} (-\bar{c} - c_{\alpha/2}, \bar{c} + c_{\alpha/2}), & 0 \leq y < \bar{c} \\ (0, y + c_{\alpha/2}), & \bar{c} \leq y < \tilde{c} \\ (g^{-1}(y), y + c_{\alpha/2}), & \tilde{c} \leq y \leq g(\bar{c} + c_{\alpha/2}) \\ (\bar{c} + c_{\alpha/2}, y + c_{\alpha/2}), & g(\bar{c} + c_{\alpha/2}) < y < \bar{c} + 2c_{\alpha/2} \\ (y - c_{\alpha/2}, y + c_{\alpha/2}), & \bar{c} + 2c_{\alpha/2} \leq y \end{cases}$$

with $\mathcal{C}(-y; \alpha) = -\mathcal{C}(y; \alpha)$. In (4.2) $g^{-1}(t)$ is well defined since g is strictly increasing to ∞ on $-\bar{c} + c_{\alpha/2} < t$, and in particular on $\tilde{c} < t$. The assumption that f is unimodal (and symmetric) ensures that (4.2) is indeed the convex hull of $\{\theta : y \in A_{MQC}(\theta)\}$.

REMARK. The MQC interval is scale invariant in the following sense: if $Y \sim (\theta, \sigma^2)$ and $Y' = Y/\sigma$, and $\mathcal{C}(y; \alpha)$ and $\mathcal{C}'(y'; \alpha)$ are the MQC confidence intervals (for any fixed ψ) based on Y and Y' , respectively, then $\mathcal{C}(y; \alpha) = \sigma \cdot \mathcal{C}'(y/\sigma; \alpha)$.

The QC and the MQC intervals determine the sign of θ starting at exactly the same value of $|t|$, but the latter constructs shorter intervals on a subset of $\{t : |t| > \bar{c}\}$ at the expense of wider intervals for all $|t| < \bar{c}$. On this subset, for each of the three cases above, the lower endpoint is farther away from zero; in the last two cases – that is, when $\psi_1 < \psi$ – there is a discontinuity point for the lower bound just when the confidence interval separates from zero ($t = \tilde{c}$).

The FCR-adjusted selective-SDCI procedure, equipped with any marginal confidence interval that satisfies requirements (MON 1) and (MON 2), has $\text{FCR} \leq q$. The actual FCR level, however, depends on the marginal confidence interval that is used. For the two-sided confidence interval – that is, for the BH-selected BY-adjusted procedure – Benjamini and Yekutieli [6] show that the FCR is also lower bounded by $q/2$. We show a similar result for the MQC interval when the estimators are Normally distributed.




Fig 1: MQC interval vs. the QC interval of [Benjamini et al.](#). The plot is for $\alpha = .05$ and $\psi = 0.7$. Both confidence intervals (weakly) determine the sign of θ whenever $|y| \geq \bar{c} = 1.81$. When the sign is determined, MQC bounds are farther away from zero for a range of y values which begins when the confidence interval separates from zero. The interval around zero where the MQC confidence limits are constant in y is the region where the QC interval includes both negative and positive values.

THEOREM 2. *For independent, Normally distributed estimators with a known variance, the FCR-adjusted selective-SDCI procedure of Definition 3 using the MQC interval with $0 < \psi < 0.9$, enjoys $FCR \geq q/2$ if $0 < q < 0.25$.*

While Theorem 2 asserts that, under the stated conditions, using the MQC interval ensures $FCR \geq q/2$, it is typically close to q . Indeed, for Normal observations and for $0 < \alpha < 0.25$ and $0 < \psi < 0.9$ the probability in (A.5) of the appendix is approximately q for all θ except for a small region where it may decrease to as low as $\alpha/2$. For example, when $\alpha = .01$ and $\psi = .85$, as long as $|\theta| \notin (0, .48)$ and $|\theta| \notin (6.43, 7.4)$, the probability in (A.5) is at least 0.99α . Hence, the inequality in (A.11) can often be made much tighter and $FCR \approx q$. We emphasize that if the original QC interval is used in the selective-SDCI procedure instead of the MQC interval, the FCR may fall significantly below $q/2$, as demonstrated in the simulation of section 5.

5. Simulation study. We carried out two different simulations that demonstrate the performance of the FCR-adjusted selective-SDCI procedure using the MQC interval. In Section A.3 of the appendix we report the results of a third simulation, in which we investigated selective-SDCI procedures under dependency.

The first simulation illustrates the asymmetric shape of the MQC intervals and its increased power to classify the sign of parameters over the BH directional procedure. We took $m = 200$ parameters where $\theta_1, \dots, \theta_{160}$ were sampled from an exponential distribution with mean 0.5, and $\theta_{161}, \dots, \theta_{200}$ were sampled from a $N(3, 1)$ distribution. Each θ_i was then randomly assigned a positive or a negative sign. The independent observations are Y_1, \dots, Y_{200} with $Y_i \sim N(\theta_i, 1)$. Figure 2 shows the constructed intervals for positive θ_i when the procedure of Definition 3 is equipped with the MQC interval ($\psi = 0.85$) and applied at level $q = 0.2$. A total of 74 sign-determining CIs were constructed, 32 of them for positive observations. The number of parameters selected is almost as large as the number selected with a BH directional procedure at level $2q = 0.4$ (77) and much larger than a BH directional procedure at level $q = 0.2$ (55). Meanwhile, the MQC constructed CIs (vertical segments in the figure) are relatively short – at the most part even shorter than the symmetric FCR-adjusted confidence intervals for level- q BH-selected parameters (partly thanks to the fact that more parameters are selected). Out of the 74 constructed CIs 14 did not cover the respective parameter (6 of which for positive observations), a proportion of 0.19. The procedure using the QC interval ($\psi = 0.85$) instead of MQC, constructed the same number of intervals with a false coverage proportion of 0.15.

The second simulation compares the (actual) FCR of the MQC-equipped procedure with that of the QC-equipped procedure. We first sampled $\theta_1, \dots, \theta_{300}$ from a $N(0, 4)$ distribution. For $N = 10^4$ datasets $\mathbf{Y} = (Y_1, \dots, Y_{300})$ with $Y_i \sim N(\theta_i, 1)$, we computed the false coverage proportion (FCP, denoted by Q_{CI} in section 2) for a level $q = 0.05$ FCR-adjusted selective-SDCI procedure using the MQC interval ($\psi = 0.85$) and for the same procedure using a QC interval ($\psi = 0.85$). We used $\psi = 0.85$ for both the QC and the MQC intervals so that sign determination occurs at the same value for both intervals. The average FCP for QC was 0.018 ($\hat{SD} = 1.3 \cdot 10^{-4}$) and for MQC it was 0.048 ($\hat{SD} = 2.2 \cdot 10^{-4}$). These results confirm that, as discussed in the previous section, the FCR when using the MQC interval is often very close to q whereas it may fall below $q/2$ when using the QC interval.

6. Detecting the sign of correlations in a social neuroscience study. Tom et al. [23] carried out an experiment in an attempt to as-

sociate neural activity in the brain with behavioral “loss aversion”. Their study received high publicity, and the collected data was reanalyzed in Poldrack and Mumford [17] and in Rosenblatt and Benjamini [19]. The original data was made available through the *OpenfMRI* initiative at <https://openfmri.org/dataset/ds000005> and described in detail in the paper by Tom et al. For each of 16 subjects a behavioral loss aversion index was measured along with a neural index at each brain voxel. The voxel-specific correlations between behavioral index and neural index were then used to detect brain regions that are associated with loss aversion.

Rosenblatt and Benjamini [19] revisited this dataset and explored different methods to construct confidence intervals which account for selection bias in reported voxels. Their approach is, in general, to employ a two-stage procedure where the first stage is in principle designed to detect nonzero correlations; at the second stage they construct a confidence interval for each parameter selected (rejected) at the first stage, while attempting to control the FCR below some pre-specified level. Specifically, one of the schemes they used is selection via the BH procedure. As we are interested in sign classification rather than two-sided testing, we view the BH procedure here as a directional procedure, namely, as a procedure which classifies the sign of each reported parameter as strictly positive or strictly negative. If willing to settle for weak (rather than strict) sign determination, our method suggests an alternative which tends to discover more parameters. Thus, we apply our method to the z -scores computed for each voxel for the Fisher-transformed correlations, and which were processed by Rosenblatt and Benjamini [19] and kindly made available to us. The concern about validity of our procedure under dependency, which is likely to be present in the current example, is mitigated by the simulations results from Section A.3.

A level 0.1 directional-BH procedure applied to the two-sided p -values found 18,844 voxels for which a strict sign decision can be made. Meanwhile, a level 0.1 FCR-adjusted selective-SDCI procedure using the MQC interval with $\psi = 0.85$ was able to weakly classify the sign of a total of 36,131 correlations, where for 27,117 of these a strict sign classification was made. For comparison, the FCR-adjusted selective-SDCI procedure using a one-sided (or Pratt’s) confidence interval, which selects according to BH procedure at level 0.2, reports 43,804 parameters, all signs weakly classified. Hence the BH at half the level makes 57% less discoveries, all with strict sign classification; whereas the MQC-equipped FCR-adjusted selective-SDCI at half the level makes only 18% less discoveries, the majority of them with strict sign classification. Figure 3a displays the MQC confidence intervals constructed for the 33,856 correlations classified as positive, along with the

QC intervals. The symmetric intervals corresponding to selection according to a level 0.1 BH procedure is also shown for reference. It is seen in the figure that for a majority of the discoveries, the lower endpoint of the MQC interval is farther away from zero than that of the QC interval, even though the latter yields the same set of discoveries. Note that the gap between the lower endpoint of MQC (black points in figure) and the lower endpoint of QC (gray line in figure) is largest immediately as the two intervals separate from the horizontal axis, which is exactly where we would like the gap to be largest: it is more important to be able to quote an endpoint farther from zero for a small detected correlation than it is for a very large detected correlation.




Fig 2: MQC sign-determining confidence intervals. Level $q = 0.2$ FCR-adjusted sign-determining MQC confidence intervals were constructed for a total of 74 out of $m = 300$ parameters, 14 of the confidence intervals do not cover the respective parameter. Vertical lines display MQC adjusted confidence intervals for the 32 positive observations: 26 of them cover the respective parameter (blue bars) and 6 of them do not (red bars). Diagonal line running through the origin is the identity line. Markers denote pairs (Y_i, θ_i) . Solid gray lines mark the position of level 0.2 FCR-adjusted two-sided CIs for level 0.2 BH-selected parameters. Broken line represents the lower boundary of the constructed QC intervals.

We emphasize that the intervals constructed by the selective-SDCI procedure using any of the configurations (i.e., any of the marginal confidence intervals) above are sign-determining. Hence, this is a partial response to the request of Rosenblatt and Benjamini [19], who comment that “it might be of interest to develop CIs that are dual to the selection methods used in neuroimaging”.




Fig 3: Selective CIs which determine the sign (left panel) and selective CIs which detect large correlations (right panel), data from Tom et al. [23]. CIs shown as vertical bars on right panel, only lower and upper endpoints of CI shown on left panel. In both panels observed correlations are on horizontal axis, vertical axis represents true correlation values; light gray solid line is the identity line. (a) Black points correspond to MQC confidence intervals, and gray lines to QC confidence intervals, for the 33,856 correlations classified as positive. The upper endpoints for the two methods coincide, while the lower endpoint of MQC is farther away from zero. (b) Requiring selective CIs to include only correlation values $\rho > 0.2$ or only values $\rho < -0.2$, the MQC_δ -equipped procedure constructs such intervals for 9 out of the original 382,362 voxels. All 9 reported confidence intervals lie above 0.2 (no correlations < -0.2 were detected).

Instead of detecting positive or negative correlations, it is reasonable that

a researcher would be interested in detecting the large correlations, positive or negative. In Section A.4 of the appendix we present an extension of the FCR-adjusted selective-SDCI procedure which allows to detect correlations $\rho_i > \rho_0$ or $\rho_i < -\rho_0$ for some pre-specified constant $\rho_0 \in (0, 1)$, and supplement decisions with compatible confidence intervals. Figure 3b displays the constructed FCR-adjusted selective confidence intervals for $\rho_0 = 0.2$.

7. Determining the direction and the type of the association in genomic association studies. In Section A.1 we offer Theorem 3 and extend the methodology developed in previous sections from the one-dimensional case to the case that $\theta_i \in \mathbb{R}^k$ with $k > 1$. We now apply the extended procedures to two-dimensional in a genomic example. The data is taken from the WTCC1 case-control study of Type-2 Diabetes [9]. We analyze data for $m = 459,653$ SNPs. The data for SNP $s = 1, \dots, m$, is a 2-by-3 table listing the SNPs' minor allele counts for the cases and for the controls. We denote by $n_{s,i,j}$ the number of subjects of type i (*controls* = 1, *cases* = 2) with $j = 0, 1, 2$ copies of a , the minor allele of SNP s . For example, the data for SNP $s = 1412$ is shown in Table 1.

TABLE 1
Data for SNP 1412.

Genotype:	AA	Aa	aa
Controls	690	1442	804
Cases	377	989	555

Clarke et al. [10] suggest using the Cochran-Armitage trend test for discovering association between SNP and disease, with weights $w = (w_1, w_2, w_3)$ that are chosen to detect particular types of association. $w = (0, 1, 1)$ is used to test whether allele a is dominant over allele A and $w = (0, 0, 1)$ is used to test whether allele a is recessive to allele A. However, most often, $w = (0, 1, 2)$ is used to test for an additive effect of allele a. The Cochran-Armitage statistic has a chi-squared distribution with 1 d.f. under the null hypothesis of no association. It is equivalent to the score statistic for the corresponding linear logit model, and its result is very similar to the logistic regression Wald test [1]. Assume that $n_{s,i,j}$ are multinomial with probabilities $(\pi_{s,1,0}, \dots, \pi_{s,2,3})$. Let $\gamma_s^j = \log(\pi_{s,2,j}/\pi_{s,1,j})$ denote the log-Odds for Diabetes for allele a count j of SNP s . For SNP 1412, the Cochran-Armitage test with $w = (0, 1, 2)$ yielded $Z = 2.605$; while the Wald test for the the null hypothesis that $\beta_s^1 = 0$ in the corresponding linear model,

$$(7.1) \quad \gamma_s^j = \beta_s^0 + \beta_s^1 \cdot j,$$

yielded $Z = 2.604$.

7.1. *Confidence regions for the Dominance and Recessiveness effects.* In our analysis we treat SNP association as a bivariate problem that corresponds to the linear model:

$$(7.2) \quad \gamma_s^j = \gamma_s^0 + \beta_s^{Dom} \cdot I(1 \leq j) + \beta_s^{Rec} \cdot I(2 \leq j).$$

Our parameters of interest are the allele a dominance effect $\beta_s^{Dom} = \gamma_s^1 - \gamma_s^0$, and the allele a recessiveness effect $\beta_s^{Rec} = \gamma_s^2 - \gamma_s^1$. For our analysis we assume that the allele effect is monotone increasing $\gamma_s^0 \leq \gamma_s^1 \leq \gamma_s^2$ or decreasing $\gamma_s^0 \geq \gamma_s^1 \geq \gamma_s^2$. Thus $(\beta_s^{Dom}, \beta_s^{Rec})$ are both either nonnegative or non-positive. In our analysis we construct confidence regions that determine the sign of $(\beta_s^{Dom}, \beta_s^{Rec})$ and indicate whether this effect is dominant ($\beta_s^{Rec} = 0$), additive ($\beta_s^{Rec} = \beta_s^{Dom}$), or recessive ($\beta_s^{Dom} = 0$).

Our parameter estimators are $\hat{\beta}_s^{Dom} = \hat{\gamma}_s^1 - \hat{\gamma}_s^0$ and $\hat{\beta}_s^{Rec} = \hat{\gamma}_s^2 - \hat{\gamma}_s^1$, with $\hat{\gamma}_s^j = \log(n_{s2j}/n_{s1j})$. To construct the confidence sets for $(\beta_s^{Dom}, \beta_s^{Rec})$, we assume that $(\hat{\beta}_s^{Dom}, \hat{\beta}_s^{Rec})$ are bivariate Normal with mean $(\beta_s^{Dom}, \beta_s^{Rec})$ and covariance matrix whose entries are the following estimated variances and covariance: $\text{Var}(\hat{\beta}_s^{Dom}) = 1/n_{s20} + 1/n_{s10} + 1/n_{s21} + 1/n_{s11}$, $\text{Var}(\hat{\beta}_s^{Rec}) = 1/n_{s21} + 1/n_{s11} + 1/n_{s22} + 1/n_{s12}$, and $\text{Cov}(\hat{\beta}_s^{Dom}, \hat{\beta}_s^{Rec}) = -1/n_{s21} - 1/n_{s11}$. Note that these parameter estimates are the same as those produced by fitting model (7.2) in R. The effect estimates for SNP 1412 are $\hat{\beta}_{1412}^{Dom} = 0.227$ and $\hat{\beta}_{1412}^{Rec} = 0.006$ and the estimated covariance matrix is

$$\hat{\Sigma}_{1412} = \begin{pmatrix} 0.0058 & -0.0017 \\ -0.0017 & 0.0047 \end{pmatrix}.$$

The confidence regions for $(\beta_{1412}^{Dom}, \beta_{1412}^{Rec})$, shown in Figure 4a, are valid under the assumption that $(\hat{\beta}_{1412}^{Dom}, \hat{\beta}_{1412}^{Rec})$ is bivariate Normal with mean $(\beta_{1412}^{Dom}, \beta_{1412}^{Rec})$ and covariance $\hat{\Sigma}_{1412}$. The black curves are equi-density curves that produce $1 - \alpha$ confidence regions with smallest volume for $(\beta_{1412}^{Dom}, \beta_{1412}^{Rec})$. The blue arrows are drawn in the direction of the principal components of the covariance matrix and their length is proportional to the square root of their variance. For $\hat{\Sigma}_{1412}$, $PC_{1412}^1 = (-0.805, 0.593)^T$ with variance 0.0071, $PC_{1412}^2 = (0.593, 0.805)^T$ with variance 0.0035.

Per construction, for all s , $\hat{\beta}_s^{Dom}$ and $\hat{\beta}_s^{Rec}$ are negatively correlated. Therefore, the 1st principal component will be a weighted difference between β_s^{Dom} and β_s^{Rec} and the 2nd principal component will be a weighted sum of β_s^{Dom} and β_s^{Rec} . As the sign of PC_s^2 is the same as the signs of β_s^{Dom} and β_s^{Rec} , we use the linear combination of $(\hat{\beta}_s^{Dom}, \hat{\beta}_s^{Rec})$ in the direction of PC_s^2 which has the smallest variance of all linear combinations of $(\hat{\beta}_s^{Dom}, \hat{\beta}_s^{Rec})$ to

determine the direction of association. The line passing through $(0, 0)$ that is perpendicular to PC_s^2 (for SNP 1412 it is the red diagonal line in Figure 1) represents 0 association. We quantify the size of association with $Z(PC_s^2)$, the distance in PC_s^2 standard deviations between $(\hat{\beta}_s^{Dom}, \hat{\beta}_s^{Rec})$ and the red diagonal. For SNP 1412, $Z(PC_{1412}^2) = 2.37$.

The $1 - \alpha$ confidence regions for $(\beta_s^{Dom}, \beta_s^{Rec})$ we propose for this problem are rectangular regions formed by intersecting a marginal $1 - \alpha_1$ confidence region for $(\beta_s^{Dom}, \beta_s^{Rec})$ in the PC_s^1 direction with a $1 - \alpha_2$ confidence region for $(\beta_s^{Dom}, \beta_s^{Rec})$ in the PC_s^2 direction, with $1 - \alpha = (1 - \alpha_1) \cdot (1 - \alpha_2)$. Orthogonality of the estimators of the principal components ensures $1 - \alpha$ coverage probability for our rectangular intervals, thereby allowing to allocate a different degree of confidence in each direction. In the next section we use this property for constructing confidence regions that are inflated due to selection *only in the direction of PC_s^2* . Furthermore, replacing the ellipsoid confidence regions with rectangular confidence regions may also lead to sharper sign determination.

The green rectangle in Figure 4a is 0.95 confidence region for $(\beta_{1412}^{Dom}, \beta_{1412}^{Rec})$ formed by intersecting a symmetric two-sided marginal 0.96 confidence interval in the PC_{1412}^2 direction and a symmetric two-sided marginal $1 - 0.0104$ confidence interval in the PC_{1412}^1 direction. Indeed, we see that since $Z(PC_{1412}^2) > z_{.02}$ this confidence region SNP is above and to the right of the red line even though the 0.95 ellipsoid bivariate Normal confidence set crosses the red diagonal, indicating positive association of allele a with Diabetes. Furthermore, the confidence region prominently covers $\beta_{1412}^{Rec} = 0$ parameter points while barely covering $\beta_{1412}^{Dom} = 0$ parameter points, suggesting that the effect of allele a is dominant.

7.2. Using FCR-adjusted CIs for determining SNP direction of association. We begin by constructing 1-dimensional confidence intervals for β_s^1 , the logistic regression coefficient for the number of minor SNP alleles in model (7.1), for all m SNPs and use these confidence intervals to determine the SNPs' direction of association.

Applying the BH procedure to $p_s = 2 \cdot (1 - \Phi(|Z(\hat{\beta}_s^1)|))$ at level $q = 0.05$ yielded 27 discoveries and at level $q = 0.10$ it yielded 43 discoveries. Here $Z(\hat{\beta}_s^1) = \hat{\beta}_s^1 / \hat{sd}(\hat{\beta}_s^1)$. Level $q = 0.05$ FCR-adjusted MQC CI with $\psi = 0.7$ yielded 35 discoveries and setting $\psi = 0.9$ yielded 36 discoveries.

We now consider the rectangular confidence regions for $(\beta_s^{Dom}, \beta_s^{Rec})$ defined in Section 7.1 that, using the algorithm in Definition 4, will be inflated for selection according to the value $Z(PC_s^2)$, for determining the SNPs' direction of association and the type of association.

For the rectangular confidence regions for $(\beta_s^{Dom}, \beta_s^{Rec})$ it is necessary to specify q_1 and q_2 , the non-coverage level for each principal component. Allocating all the non-coverage probability to PC_s^2 , i.e. setting $q_2 = q$ and $q_1 = 1$, reduces the confidence regions to $1 - q$ marginal confidence intervals for PC_s^2 and level q selection rules that are based on PC_s^2 . Applying the BH procedure to $p_s = 2 \cdot (1 - \Phi(|Z(PC_s^2)|))$ at level $q = 0.05$ yielded 23 discoveries and at level $q = 0.10$ it yielded 31 discoveries. Level $q = 0.05$ FCR-adjusted MQC CI with $\psi = 0.7$ yielded 24 discoveries and setting $\psi = 0.9$ yielded 30 discoveries.

We now consider confidence regions for $(\beta_s^{Dom}, \beta_s^{Rec})$ with $q_2 = 0.04$ and $q_1 = 0.0104$. As selection is only applied in the direction of PC_s^2 , the CIs for PC_s^2 are level $q_2 = 0.04$ FCR-adjusted marginal CIs that are based on $Z(PC_s^2)$, and the CIs for PC_s^1 are (unadjusted) $1 - 0.0104$ marginal two-sided CIs based on $Z(PC_s^1)$. Applying the BH procedure to $P_s = 2 \cdot (1 - \Phi(|Z(PC_s^2)|))$ at level $q = 0.04$ yielded 23 discoveries as before, and at level $q = 0.08$ the BH procedure yielded 30 discoveries. Level $q = 0.04$ FCR-adjusted MQC CI with $\psi = 0.7$ yielded 24 discoveries, and setting $\psi = 0.9$ yielded 29 discoveries. Even though the distribution of $|Z(\hat{\beta}_s^1)|$ was larger than that of $|Z(PC_s^2)|$ (it yielded more discoveries) the ordering of the SNPs according to the two Z-scores was very similar – the ranking of the 19 most significant SNPs according to $Z(\hat{\beta}_s^1)$ and $Z(PC_s^2)$ was the same.

SNP 69962 has large $Z(PC_{69962}^2) = 4.461$ (ranked 27) and a relatively smaller $Z(\hat{\beta}_{69962}^1) = 4.074$ (ranked 64 – undiscoverable with level $q = 0.10$ BH procedure). The rectangle formed by the solid green lines in Figure 4b is the MQC CI rectangular confidence sets for $(\beta_{69962}^{Dom}, \beta_{69962}^{Rec})$ with $q_2 = 0.04$ and $q_1 = 0.0104$ and $\psi = 0.9$ adjusted for the 29 selected SNPs. The smaller rectangle formed by the broken green lines and the solid green lines is the unadjusted confidence sets for $(\beta_{69962}^{Dom}, \beta_{69962}^{Rec})$ with $q_2 = 0.04$ and $q_1 = 0.0104$. As SNP 69962 was selected its selection adjusted confidence set does not cross the red line, but rather is on the red line corresponding to MQC CI lower boundary that is equal 0; this indicates non-negative association with Diabetes. Furthermore, the fact that the selective confidence region prominently covers $\beta_{69962}^{Rec} = 0$ parameter points while barely covering $\beta_{69962}^{Dom} = 0$ parameter points suggests that the effect of allele a is dominant.

8. Discussion. Selective inference refers to the general situation where the target of inference is chosen adaptively – only after seeing the data. We concentrated on a setup where selective inference arises in connection to multiplicity: the analyst collects noisy observations on a (typically large) number m of unknown parameters, which he will use to first try and answer

a primary question about each parameter, and second to construct CIs for only the parameters for which there was enough evidence to answer the primary question. Specifically, we considered the problem of detecting the sign of parameters, and supplementing each directional decision made with a CI. Because the same data is used for detection and for construction of the follow-up CIs, selection need be accounted for.

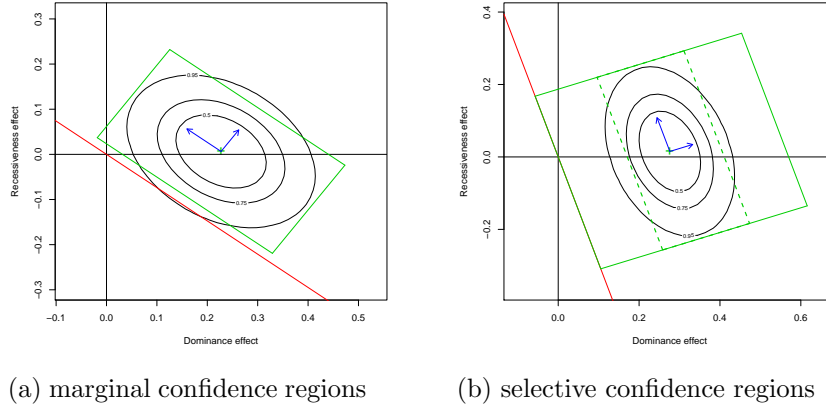


Fig 4: (a) Confidence regions for $(\beta_{1412}^{Dom}, \beta_{1412}^{Rec})$. Green plus sign is $(\hat{\beta}_{1412}^{Dom}, \hat{\beta}_{1412}^{Rec})$. Black curves are 0.50, 0.75 and 0.95 bivariate Normal confidence regions. Blue arrows are the principal components of the bivariate Normal distribution. Green rectangle is a 0.95 confidence set based on the two principal components. Red line is drawn at association effect equals 0. (b) Selection-adjusted confidence regions for $(\beta_{69962}^{Dom}, \beta_{69962}^{Rec})$. Green plus sign is at $(\hat{\beta}_{69962}^{Dom}, \hat{\beta}_{69962}^{Rec})$. Black curves are 0.50, 0.75 and 0.95 (unadjusted) bivariate Normal confidence regions. Green rectangle captured by broken lines is the 0.95 unadjusted confidence set based on the two principal components. Solid green rectangle is the 0.95 selection-adjusted confidence set based on the two principal components. Red line is drawn at association effect equals 0.

Besides validity, we emphasized the compatibility of a constructed CI with the directional decision. To ensure compatibility, instead of the natural two-stage approach that treats selection (detection) and CI construction separately, we suggested a procedure that is entirely driven by CIs. Thus, the task becomes that of constructing selective CIs which include either only positive values or only non-positive values, while controlling the FCR; as the selection criterion is determined by the CIs themselves, compatibility is

guaranteed per construction. Constructing CIs that are valid *conditional* on selection, on the other hand, is almost by definition destined to fail when it comes to ensuring compatibility; the problem is essentially that, paraphrasing the words of Fithian et al. [12], conditioning precludes us from appealing to the evidence in the data that allowed us to make a decision on the sign in the first place.

While the focus was predominantly on the sign problem, the approach we suggest is quite general. In Section A.4, for example, we showed how to modify the procedure of Definition 3 so that instead of sign classification, the primary goal is to detect parameters larger than δ or smaller than $-\delta$. In general, suppose that the data is $Y_i \stackrel{\text{ind}}{\sim} f(y; \theta_i)$, $i = 1, \dots, m$ where $Y_i \in \mathcal{Y}$ and $\theta_i \in \Theta$. Let $\Theta_j \subseteq \Theta$, $j = 1, \dots, k$ be disjoint subsets in the parameter space. The primary task is to detect membership of the θ_i to any of the Θ_j ; the secondary task is to construct a confidence set C_i for each classified parameter, such that $C_i \subseteq \Theta_j$ if θ_i was classified to Θ_j . For hypothesis testing $k = 1$ and Θ_1 is the set of alternatives; for the (weak) sign problem $k = 2$ and $\Theta_1 = (0, \infty)$, $\Theta_2 = (-\infty, 0]$; the example of Section A.4 corresponds to $k = 2$ and $\Theta_1 = (\delta, \infty)$, $\Theta_2 = (-\infty, -\delta)$; in Section 7 $\mathcal{Y} = \Theta = \mathbb{R}^2$ and $k = 2$ and $\Theta_1 = (-\infty, \infty) \times (0, \infty)$, $\Theta_2 = (-\infty, \infty) \times (-\infty, 0]$. In principle, the extension of the procedure in Definition 3 to the general case would be to construct the maximum number of FCR-adjusted confidence sets such that each confidence set is contained in one of the subsets Θ_j , $j = 1, \dots, k$.

There are certainly remaining challenges. When $Y_i \sim N(\theta_i, \sigma^2)$ and σ is unknown, Finner [11, Section 3] pointed out that a disadvantage of the the CIs based on the t statistic is that they are unbiased for the (natural) parameter θ_i/σ , not θ_i , and suggested an alternative CI which improves uniformly over the t procedure. It might be of interest to try and modify Finner's CI to produce an interval with similar properties as the MQC interval; for the corresponding procedure of Definition 3 to be valid, the monotonicity requirements would need to be checked, which might not be trivial. Another direction worth exploring is constructing sign-determining CIs for coefficients β_j in a linear regression model; Barber and Candès [2] address sign classification under directional-FDR control in the Gaussian linear model. To supplement such directional decisions with compatible confidence bounds is of clear practical importance.

Another important issue is establishing a benchmark against which our procedure can be evaluated: while our procedure balances between power and length of constructed CIs, it is indexed by a single scalar parameter (ψ); it is natural to ask if more can be gained – for example, in the form of shorter CIs – when allowing more flexibility in constructing selective CIs that

determine the sign. To be able to compare different procedures, a reasonable option is to set up a formal criterion which will take into account both power and the shape of constructed intervals.

APPENDIX A: APPENDIX SECTION

A.1. Sign determination by confidence regions. In this section we extend our methodology from the one-dimensional case to the case that $\theta_i \in \mathbb{R}^k$ with $k > 1$. In principle, it is possible to classify θ_i as having one of 2^k possible signs. In this paper we consider a straightforward extension of the 1-dimensional case in which $\theta_i = (\theta_{i1}, \theta_{i2})$, with $\theta_{i1} \in \mathbb{R}^{k-1}$ and $\theta_{i2} \in \mathbb{R}^1$ where we try to classify θ_i according to the sign of θ_{i2} . In the next section we apply this methodology for classifying the sign of the genetic association of almost half a million SNPs and then assessing whether the effect of the SNP is recessive or dominant.

Let $Y_i = (Y_{i1}, Y_{i2})$, with independent $Y_{i1} \sim f(y_{i1} - \theta_{i1})$ and $Y_{i2} \sim f(y_{i2} - \theta_{i2})$. $CI_{i1}(\alpha) = CI_{i1}(\alpha; Y_{i1})$ is a marginal $1 - \alpha$ confidence interval for θ_{i1} and $CI_{i2}(\alpha) = CI_{i2}(\alpha; Y_{i2})$ is a marginal $1 - \alpha$ confidence interval for θ_{i2} (assume that the coverage probability is exactly $1 - \alpha$, not more). We use $CI_{i1}(\alpha_1)$ and $CI_{i2}(\alpha_2)$ to form a $1 - \alpha_1 \cdot \alpha_2$ confidence set for θ_i ,

$$\widetilde{CI}_i(\alpha_1, \alpha_2) = \{\theta_i : \theta_{i1} \in CI_{i1}(\alpha_1), \theta_{i2} \in CI_{i2}(\alpha_2)\}.$$

For independent Y_1, \dots, Y_m and a selection rule that has $R_{\min}(\mathbf{Y}^{(i)}) \equiv R_{CI}$, BY show that the FCR is equal to $\sum_{r=1}^m \sum_{i=1}^m \Pr(|S| = r, \widetilde{NCI}_i) / r$, where \widetilde{NCI}_i is the event that θ_i is selected and $\theta_i \notin \widetilde{CI}_i$. We denote by \widetilde{NCI}_{i1} the event that θ_i is selected and $\theta_{i1} \notin CI_{i1}$ and by \widetilde{NCI}_{i2} the event that θ_i is selected and $\theta_{i2} \notin CI_{i2}$. Thus $\widetilde{NCI}_i = \widetilde{NCI}_{i1} \cup \widetilde{NCI}_{i2}$, and to evaluate FCR we express \widetilde{NCI}_i as the disjoint union

$$\widetilde{NCI}_i = \widetilde{NCI}_{i1} \cup (\{\theta_{i1} \in CI_{i1}\} \cap \widetilde{NCI}_{i2}).$$

We consider selection rules $\mathcal{S}_2(\mathbf{Y}_{\bullet 2})$ that are determined by only $\mathbf{Y}_{\bullet 2} = (Y_{12}, \dots, Y_{m2})$.

DEFINITION 4. Level- (q_1, q_2) FCR-Adjustment for Selection Rules Determined by $\mathbf{Y}_{\bullet 2}$

1. Apply the selection criterion \mathcal{S}_2 to obtain $\mathcal{S}_2(\mathbf{Y}_{\bullet 2})$.
2. For each selected parameter θ_i , $i \in \mathcal{S}_2(\mathbf{Y}_{\bullet 2})$, let

$$(A.1) \quad R_{\min}(\mathbf{Y}_{\bullet 2}^{(i)}) = \min_t \left\{ \left| \mathcal{S}(\mathbf{Y}_{\bullet 2}^{(i)}, Y_{i2} = t) \right| : i \in \mathcal{S}(\mathbf{Y}_{\bullet 2}^{(i)}, Y_{i2} = t) \right\},$$

where $\mathbf{Y}_{\bullet 2}^{(i)}$ is the vector obtained by omitting Y_{i2} from $\mathbf{Y}_{\bullet 2}$.

3. For each selected parameter θ_i , $i \in \mathcal{S}_2(\mathbf{Y}_{\bullet 2})$, construct the following CI:

$$\widetilde{CI}_i \left(q_1, \frac{R_{\min}(\mathbf{Y}_{\bullet 2}^{(i)}) \cdot q_2}{m} \right).$$

THEOREM 3. *Let $\mathbf{Y}_1, \dots, \mathbf{Y}_m$ be independent where $\mathbf{Y}_i = (Y_{i1}, Y_{i2})$ for independent Y_{i1} and Y_{i2} . Then the FCR of the level- (q_1, q_2) FCR-adjustment for $\mathcal{S}_2(\mathbf{Y}_{\bullet 2})$ is*

$$FCR(\widetilde{CI}_{\bullet}; \mathcal{S}; q_1, q_2) = q_1 + (1 - q_1) \cdot FCR(CI_{\bullet 2}; \mathcal{S}_2; q_2)$$

where $FCR(CI_{\bullet 2}; \mathcal{S}_2; q_2)$ is the FCR of the level q_2 BY FCR-adjusted CI for \mathcal{S}_2 .

A.2. Connections to Existing Work on Post-Selection Inference.

In the procedure of Definition 3 a confidence interval is constructed for θ_i only if $i \in \mathcal{S}(\mathbf{Y})$ where \mathcal{S} is a pre-specified rule. Viewed as a two-stage procedure (as defined in Section 1), the second stage must take selection into account in order to maintain FCR control, but adjusting the level of marginal confidence intervals is not the only way to achieve it. To restore validity of post-selection confidence intervals, a common approach, which is not limited to the multiplicity setup or to selection rules we consider in the current paper, is to construct intervals that have the nominal coverage level *conditional* on selection. Thus, in the setup of the current paper, whenever $i \in \mathcal{S}(\mathbf{Y})$ the conditional approach would construct an interval $CI_i(q)$ with the property

$$(A.2) \quad \Pr(\theta_i \in CI_i(q) | i \in \mathcal{S}(\mathbf{Y})) \geq 1 - q.$$

Conditional CIs based on a truncated univariate Normal distribution were suggested in Zhong and Prentice [25] and Weinstein et al. [24]. A Recent line of work, including Lee et al. [14], Taylor et al. [21], Sun and Taylor [20] (among others), greatly increased the applicability of the conditional approach by developing exact confidence intervals methods when selection corresponds to truncating a multivariate Normal distribution to a polyhedron; these results were in turn extended to generalized linear models in Fithian et al. [12], who also provided further theoretical support.

Under independence of the Y_i , we obtain a valid two-stage procedure by selecting parameters through a level- $2q$ BH procedure for testing $H_0^i : \theta_i = 0$, classifying θ_i as positive or non-positive according as $Y_i > 0$ or $Y_i < 0$; then for each classified parameter constructing a $1 - q$ conditional

CI. From Weinstein et al. [24, Section 7] it follows that if for each $i \in \mathcal{S}(\mathbf{Y})$ a confidence interval $CI_i = \mathcal{C}(Y_i; \sigma^2, \hat{c}, q)$ is constructed where $\mathcal{C}(Y_i; \sigma^2, c, \alpha)$ has the property that

$$(A.3) \quad \Pr_{Y \sim N(\theta, \sigma^2)} (\theta \in \mathcal{C}(Y; \sigma^2, c, \alpha) \mid |Y| > c) \geq 1 - \alpha$$

and where

$$\hat{c} = \Phi^{-1}(1 - i^*q/m), \quad i^* = \max\{i : P_{(i)} \leq i(2q)/m\}$$

for $P_i = 2(1 - \Phi(|Y_i|))$ and $P_{(1)} \leq P_{(2)} \leq \dots P_{(m)}$, then (A.2) is satisfied. This is because the conditional distribution of $Y_i \mid (\mathbf{Y}^{(i)}, i \in \mathcal{S}(\mathbf{Y}))$ is that of a Normal variable truncated to $\{y : |y| > \hat{c}\}$ where \hat{c} is a constant determined by $\mathbf{Y}^{(i)}$.

The procedure described above controls both $\text{wdFDR} \leq q$ and $\text{FCR} \leq q$ as required; however, a disadvantage is that – unlike the procedure of Definition 3 – it cannot ensure that a constructed CI_i does not cross zero; this would contradict with the fact that the sign of θ_i was detected. In the example of Section 6, if we use the conditional CI of Weinstein et al. [24, CQC with $\lambda = 0.4$] after BH selection at level 0.2, then 28,082 out of the 43,804 constructed intervals include both positive and negative values. Hence 64% of the parameters whose sign was classified at the first stage, are supplemented with intervals that do not determine the sign. In fact, the construction of Weinstein et al. [24] is designed specifically to promote sign detection, and still it is not able to guarantee it.

We now show that the situation would be similar if any other conditional CI is used instead. Indeed, for a constant c let $\bar{P}_\theta(\cdot)$ indicate probability under the conditional distribution of $Y \sim N(\theta, 1)$ given $|Y| > c$, and suppose that there exists a procedure $\mathcal{C}(y; \alpha) = \mathcal{C}(y; \sigma^2 = 1, c, \alpha)$ with the property (A.3) such that $\mathcal{C}(y; \alpha) \subseteq (-\infty, 0]$ or $\mathcal{C}(y; \alpha) \subseteq (0, \infty)$ for all $|y| > c$. Let $\alpha < 1/2$. Then

$$\begin{aligned} 1 &= \bar{P}_0(\mathcal{C}(Y; \alpha) \subseteq (-\infty, 0]) + \bar{P}_0(\mathcal{C}(Y; \alpha) \subseteq (0, \infty)) \\ &= \lim_{\theta \rightarrow 0^+} \bar{P}_\theta(\mathcal{C}(Y; \alpha) \subseteq (-\infty, 0]) + \bar{P}_0(\mathcal{C}(Y; \alpha) \subseteq (0, \infty)) \leq \alpha + \alpha = 2\alpha \end{aligned}$$

which is a contradiction. The equality in the second line is by continuity of $\bar{P}_\theta(\mathcal{C}(y; \alpha) \subseteq (-\infty, 0])$ at $\theta = 0$. The inequality in the third line is because for any $\theta > 0$, coverage property of the interval implies necessarily that $\bar{P}_\theta(\mathcal{C}(y; \alpha) \subseteq (-\infty, 0]) \leq \alpha$, and similarly $\bar{P}_0(\mathcal{C}(y; \alpha) \subseteq (0, \infty)) \leq \alpha$.

A.3. Selective-SDCI Procedures Under Dependency. Equipped with any marginal CI that satisfies the requirements (MON 1) and (MON2),

a level- q selective-SDCI procedure is guaranteed to control $\text{FCR} \leq q$ under independence of the observations. The case of dependent data is considerably more challenging. For general dependency, applying the selective-SDCI of Definition 3 (equipped with any marginal CI) at level $q / \sum_{j=1}^m \frac{1}{j}$ ensures $\text{FCR} \leq q$. This follows immediately from Theorem 4 in Benjamini and Yekutieli [6].

If the estimators are positive regression dependent on a subset [PRDS hereafter; 5], a consequence of Theorem 3 in Benjamini and Yekutieli [6] is that if the level- q procedure of Definition 3 is equipped with a the interval $\mathcal{C}(y; \alpha) = (y - c_\alpha, \infty)$ (or $\mathcal{C}(y; \alpha) = (-\infty, y + c_\alpha)$), then the FCR is still controlled at q under PRDS; however, if sign detection is of interest, we would never want to equip the selective-SDCI with such a CI (which would mean giving up on either detection of positive or on detection of negative parameters). Hence Theorem 3 in Benjamini and Yekutieli [6] does not really cover selection followed by confidence interval construction via the selective-SDCI procedure of Definition 3. We would like to point out that under PRDS even the validity of the level- q directional-BH procedure, which is a special case of our selective-SDCI procedure, has not been established before.

We examine FCR of the selective-SDCI procedure in practice using simulations. The data simulates the brain voxel data of section 6 under dependency; it was generated as in Rosenblatt and Benjamini [19] and using their code, available at <https://github.com/johnros/SelectiveEstimationSimulations>. Specifically, for each configuration of nonzero effect size and proportion of nulls, and in each of 100 rounds, data representing z -transformed ($n = 16$ subjects) correlations for a “brain” of $10 \times 10 \times 10$ voxels was generated as the sum of a signal field and a smooth Gaussian noise field. The smoothness of the noise field – controlling the spatial covariance – is represented by the parameter FWHM, which was varied at 3 different levels $\{3.3, 4.7, 5.7\}$ (4.7 being the estimated quantity from the actual data analyzed in section 6). The smooth Gaussian random field used in generating the data is PRDS [16], hence is an appropriate case to examine the actual FCR of our procedure for PRDS data. More details describing how data was generated are available in Rosenblatt and Benjamini [19, Appendix C.2].

In each round we applied our procedure at level $q = 0.1$, first using the QC interval as the marginal CI, and second using the MQC as the marginal CI. For each of the two methods we recorded the proportion of non-covering constructed CIs (FCP) as well as the number of constructed sign-determining intervals. The results are presented in Figure 5, which shows that, overall, the situation is qualitatively similar to the independent case: for almost each

simulation configuration the estimated FCR is under 0.1, while it is much closer to the nominal level for the MQC-equipped procedure. Specifically, for larger proportion of non-zero effects ($\pi_1 = 0.5, 0.9$), the estimated FCR of the MQC-equipped procedure is larger than $q/2 = 0.05$ for all configurations of smoothness (FWHM) and levels of non-zero signal strength (≥ 0.2).




Fig 5: Estimated FCR of the selective-SDCI procedure under dependency. Each point in the figure is estimated FCR for a specific configuration, when applying the procedure of Definition 3 at level 0.1 to Fisher-transformed voxel correlations (sample size is $n = 16$). Vertical bars are drawn at \pm two standard errors. In each simulation round the underlying signal for each voxel was independently set to ρ_1 w.p. π_1 or zero w.p. $1 - \pi_1$. The two line types correspond to the procedure using the QC marginal interval (broken) and the MQC marginal interval (solid). For the no-signal case ($\pi_1 = 0$) the two lines coincide.

A.4. Detecting large correlations. The focus in Section 6 was on detecting the sign of correlations. In practice it may be of interest to detect instead only large correlations, namely, correlations $\rho_i > \rho_0$ or $\rho_i < -\rho_0$ for some pre-specified constant $\rho_0 \in (0, 1)$, while still controlling the proportion of incorrect decisions. Hodges and Lehmann [13] referred to testing of the interval hypothesis $H_0 : \rho \in [-\rho_0, \rho_0]$ as testing for “material significance”. Also, as pointed out by Finner [11], our pursuit reflects in some sense the opposite goal of the bioequivalence problem, where the aim is to detect parameters $|\rho| < \rho_0$. We offer here an extension of our selective-SDCI

procedure and use it to detect large correlations in the study of Section 6.

As before, let $Y_i \sim f(y_i - \theta_i)$ $i = 1, \dots, m$, with f a unimodal and symmetric density. Fix $\delta \in (0, \infty)$. We are interested in a procedure which, for a subset $S = \mathcal{S}(Y) \in \{1, \dots, m\}$, constructs a confidence interval CI_i for each $\theta_i, i \in S$, such that $CI_i \subseteq (\delta, \infty)$ or $CI_i \subseteq (-\infty, -\delta)$, and at the same time $\text{FCR} \leq q$. Adapting the selective-SDCI procedure from the original sign problem to the current situation is straightforward: simply replace R in Definition 3 with

$$R = \max \left\{ r : CI_{(r)} \left(\frac{r \cdot q}{m} \right) \text{ includes only values } > \delta \text{ or only values } < -\delta \right\}$$

and leave the procedure of Definition 3 otherwise unchanged. As was the case for the sign problem, any marginal confidence interval satisfying Requirements (MON 1) and (MON 2) can be used with the selective-SDCI procedure. However, in order to obtain a powerful procedure, we would like to use a marginal confidence interval which is suited to the current task rather than to the sign problem. Specifically, we adapt the MQC marginal interval of section 4 to obtain an interval that avoids intersecting $[-\delta, \delta]$ starting at an observation value smaller (in absolute value) than $\delta + c_{\alpha/2}$. To save space, we leave the description of the adapted MQC interval to the appendix. In Figure 6 the shape of this interval, which we refer to as MQC_δ , is shown for a Normal density and $\delta = 0.5$.

We set $q = 0.1$ and applied the procedure of Definition 3 equipped with the MQC_δ interval to the data from Tom et al. [23] for detecting correlations $\rho_i > 0.2$ or $\rho_i < -0.2$; the value $\rho_0 = 0.2$ was chosen to represent a “sufficiently large” correlation size. The results are shown in figure 3b. Out of the 382,362 voxels originally considered in our analysis, our procedure finds that only 9 can be declared to have a correlation larger than 0.2 or smaller than -0.2. All 9 reported intervals are for positive estimated correlations, and cover only values larger than 0.2. Hence, while detection of the sign was possible for as many as 36,131 parameters (in the previous subsection), there is a dramatic decrease in the number of finding already when we aim at discovering correlations of size at least 0.2.

Note that the 9 constructed intervals which lie above 0.2 are FCR-adjusted at a more stringent level than the sign-determining intervals constructed before, as less parameters are selected. This inflation is needed to ensure that FCR is controlled for the new set of findings: reporting those among the 36,131 sign-determining intervals which further lie above 0.2 or below -0.2 would, of course, suffer from the same type of selection bias problem as would reporting in the first place all unadjusted 90% intervals which lie above 0.2 or below -0.2.




Fig 6: Modified Quasi-Conventional confidence interval MQC_δ for early detection of large effects. Here $Y \sim N(\theta, 1)$ and the confidence interval for θ is designed to exclude values in $[-\delta, \delta]$ as early as possible, i.e., for small values of $|y|$. The plot is for $\alpha = 0.1, \delta = 0.5$. For this configuration the interval lies completely above 0.5 or completely below -0.5 starting at $|y| = 1.84$; compare to $\delta + z_{1-0.1/2} = 2.14$ for the symmetric interval, shown in gray. Broken lines are drawn at $\pm\delta$ and the origin is marked with a plus sign.

A.5. A Full specification of MQC CI. In Section 4 (equation (4.2)) we provided a specification of the MQC confidence interval when $0 < \psi \leq \psi_1$. A complete specification of the MQC confidence interval follows. The constants $\bar{c}, \tilde{c}, \psi_1, \psi_2$ and the function $g(\theta)$ are all as defined in Section 4.

If $0 < \psi \leq \psi_1$,

$$\mathcal{C}_{MQC}(y; \alpha) = \begin{cases} (-\bar{c} - c_{\alpha/2}, \bar{c} + c_{\alpha/2}), & 0 \leq y < \bar{c} \\ [0, y + c_{\alpha/2}), & \bar{c} \leq y < c_{\alpha/2} \\ (0, y + c_{\alpha/2}), & c_{\alpha/2} \leq y < \tilde{c} \\ (g^{-1}(y), y + c_{\alpha/2}), & \tilde{c} \leq y \leq g(\bar{c} + c_{\alpha/2}) \\ (\bar{c} + c_{\alpha/2}, y + c_{\alpha/2}), & g(\bar{c} + c_{\alpha/2}) < y < \bar{c} + 2c_{\alpha/2} \\ (y - c_{\alpha/2}, y + c_{\alpha/2}), & \bar{c} + 2c_{\alpha/2} \leq y \end{cases}$$

If $\psi_1 < \psi \leq \psi_2$,

$$\mathcal{C}_{MQC}(y; \alpha) = \begin{cases} (-\bar{c} - c_{\alpha/2}, \bar{c} + c_{\alpha/2}), & 0 \leq y < \bar{c} \\ [0, y + c_{\alpha/2}), & \bar{c} \leq y < c_{\alpha/2} \\ (0, y + c_{\alpha/2}), & c_{\alpha/2} \leq y < \tilde{c} \\ (\bar{c} + c_{\alpha/2}, y + c_{\alpha/2}), & \tilde{c} \leq y \leq \bar{c} + 2c_{\alpha/2} \\ (y - c_{\alpha/2}, y + c_{\alpha/2}), & \bar{c} + 2c_{\alpha/2} < y \end{cases}$$

If $\psi_2 < \psi$,

$$\mathcal{C}_{MQC}(y; \alpha) = \begin{cases} (-\bar{c} - c_{\alpha/2}, \bar{c} + c_{\alpha/2}), & 0 \leq y < \bar{c} \\ [0, y + c_{\alpha/2}), & \bar{c} \leq y < c_{\alpha/2} \\ (0, y + c_{\alpha/2}), & c_{\alpha/2} \leq y < \tilde{c} \\ (y - c_{\alpha/2}, y + c_{\alpha/2}), & \tilde{c} \leq y \end{cases}$$

with $\mathcal{C}(-y; \alpha) = -\mathcal{C}(y; \alpha)$.

A.6. Technical Details for MQC_δ . We give the technical details for deriving MQC_δ shown in Figure 6. For a fixed $\delta > 0$, the construction presented below can be viewed as adapting the MQC interval so that instead of determining the sign, it lies above δ or below $-\delta$ for relatively small (absolute) values of the observation.

Like the MQC interval of Section 4, the new interval is obtained by inverting a family of acceptance regions. The $1 - \alpha$ acceptance regions describing the confidence interval are

$$A_{MQC_\delta}(\theta) = \begin{cases} (-\delta - \bar{c}, \delta + \bar{c}), & 0 \leq \theta \leq \delta \\ (-\delta - \bar{c}, g_\delta(\theta)), & \delta \leq \theta < \delta + \bar{c} + c_{\alpha/2} \\ (\theta - c_{\alpha/2}, \theta + c_{\alpha/2}), & \delta + \bar{c} + c_{\alpha/2} < \theta \end{cases}$$

with $A_{MQC_\delta}(\theta) = -A_{MQC_\delta}(\theta)$ for $\theta < 0$. Above, the constant \bar{c} is determined by δ through

$$(A.4) \quad F(\delta + \bar{c} - \delta) - F(-\delta - \bar{c} + \theta) = 1 - \alpha$$

and the function g_δ is given by

$$g_\delta(\theta) = \theta + F^{-1}\{1 - \alpha + F(-\delta - \bar{c} - \theta)\}.$$

The convex hull of $\{\theta : y \in A_{MQC_\delta}(\theta)\}$ is then

$$\mathcal{C}(y; \alpha) = \begin{cases} (-\delta - \bar{c} - c_{\alpha/2}, \delta + \bar{c} + c_{\alpha/2}), & 0 \leq x < \delta + \bar{c} \\ (g^{-1}(x), x + c_{\alpha/2}), & \delta + \bar{c} \leq x < g(\delta + \bar{c} + c_{\alpha/2}) \\ (\delta + \bar{c} + c_{\alpha/2}, x + c_{\alpha/2}), & g(\delta + \bar{c} + c_{\alpha/2}) \leq x < \delta + \bar{c} + 2c_{\alpha/2} \\ (x - c_{\alpha/2}, x + c_{\alpha/2}), & \delta + \bar{c} + 2c_{\alpha/2} \leq x \end{cases}$$

Figure 6 shows the resulting interval for a Normal distribution and $\delta = 0.5$. Note that while the MQC interval was parametrized by ψ (or, equivalently, by \bar{c}) which determined the tradeoff between early sign determination and maximum length of the confidence interval, MQC_δ is not indexed by such a parameter. Indeed, for MQC_δ there is no flexibility in choosing how early the interval stops crossing δ (or $-\delta$): To any δ corresponds a constant \bar{c} given by (A.4). Note also that while MQC guarantees only weak determination of the sign (≤ 0 or > 0) whenever it does not include values of both signs, MQC_δ is always an open interval, and it separates from δ ($-\delta$) immediately at $\delta + \bar{c}$, whereas MQC separates from zero at $\tilde{c} > \bar{c}$. These are consequences of the difference between the sign problem and the problem of detecting large effects.

A.7. A proof that $R_{\min}(\mathbf{Y}^{(i)}) = R_{CI}(\mathbf{Y})$ in Theorem 1. Without loss of generality, we show that $|\mathcal{S}^*(\mathbf{Y}^{(1)}, Y_1 = y)|$ is constant over y for all y is such that $i \in \mathcal{S}^*(\mathbf{Y}^{(1)}, Y_1 = y)$. Let

$$g(\alpha) = \inf\{y \geq 0 : \mathcal{C}(y; \alpha) \text{ includes values of one sign only}\},$$

and let $\tau(i) = g(\frac{i}{m}q)$, $i = 1, \dots, m$. Recall that $Y_{(i)}$ is the i -th *largest* value among $|Y_1|, \dots, |Y_m|$. Because $\mathcal{C}(y; \alpha)$ satisfies the monotonicity requirements (MON 1) and (MON 2), $i^* = \max\{i : \tau(i) \leq Y_{(i)}\}$ and $\tau(i)$ is a decreasing sequence. Define now a vector $\tilde{\mathbf{Y}} = (\tilde{\mathbf{Y}}^{(1)}, \tilde{Y}_1)$, which depends on $\mathbf{Y}^{(1)}$ only, by $\tilde{\mathbf{Y}}^{(1)} = \mathbf{Y}^{(1)}$, $\tilde{Y}_1 = \infty$. Let $\tilde{Y}_{(i)}$ denote the i -th *largest* value among $|\tilde{Y}_1|, \dots, |\tilde{Y}_m|$. Furthermore, let

$$\tilde{i}^* = \max\{1 \leq i \leq m : \tau(i) \leq \tilde{Y}_{(i)}\}.$$

We will show that if $1 \in \mathcal{S}^*(\mathbf{Y})$ then $i^* = \tilde{i}^*$, hence if $1 \in \mathcal{S}^*(\mathbf{Y})$ then $|\mathcal{S}^*(\mathbf{Y})| = i^*$, which does not depend on y .

First, note that $1 \in \mathcal{S}^*(\mathbf{Y}) \iff \tau(\tilde{i}^*) \leq |y|$. Indeed, suppose that $y < \tau(\tilde{i}^*)$; then for all $i \geq \tilde{i}^*$, $Y_{(i)} \leq \min(\tilde{Y}_{(i)}, \tau(\tilde{i}^*))$. Therefore, for all $i \geq \tilde{i}^*$, $Y_{(i)} < \tau(i)$, which together with the fact that $\tau(i)$ is decreasing implies that $i \notin \mathcal{S}^*(\mathbf{Y})$ if $|Y_i| < \tau(\tilde{i}^*)$. In particular, $1 \notin \mathcal{S}^*(\mathbf{Y})$. On the other hand, if

$\tau(\tilde{i}^*) \leq |y|$, then $Y_{(\tilde{i}^*)} = \min(\tilde{Y}_{(\tilde{i}^*)}, |y|) \geq \tau(\tilde{i}^*)$, which together with the fact that $\tau(i)$ is decreasing implies that $i \in \mathcal{S}^*(\mathbf{Y})$ if $\tau(\tilde{i}^*) \leq |Y_i|$. In particular, $1 \in \mathcal{S}^*(\mathbf{Y})$.

To complete the proof, observe that when $\tau(\tilde{i}^*) \leq |y|$, (i) $Y_{(i)} < \tau(i)$ for $i > \tilde{i}^*$, which implies $i^* \leq \tilde{i}^*$, and (ii) $Y_{(\tilde{i}^*)} = \min(\tilde{Y}_{(\tilde{i}^*)}, |y|) \geq \tau(\tilde{i}^*)$, which implies that $i^* \leq \tilde{i}^*$. We conclude that $i^* = \tilde{i}^*$, as required.

A.8. Proof of Theorem 2. By the remark in Section 4, it is enough to prove the theorem for the case $\text{Var}(Y_i) = 1$. Indeed, for $\sigma^2 = \text{Var}(Y_i)$, letting $Y'_i = Y_i/\sigma$ and $\theta'_i = \theta_i/\sigma$ we have that $\theta_i \notin \mathcal{C}(Y_i; \alpha) \iff \theta'_i \notin \mathcal{C}'(Y'_i; \alpha)$ where $\mathcal{C}(y; \alpha)$ and $\mathcal{C}'(y'; \alpha)$ are the MQC CIs corresponding to the distributions of Y and Y' , respectively. Therefore the FCR of the procedure defined for the Y_i (w.r.t. the θ_i) is the same as the procedure defined for the Y'_i (w.r.t. the θ'_i).

First we claim that for $\psi < 0.9$, the MQC interval is given by (4.2) for all $0 < \alpha < 0.25$. We need to check that $\psi_1 > 0.9$ for all $0 < \alpha < 0.25$. It can be verified that ψ_1 is a decreasing function of α on $0 < \alpha < 0.25$, and we have $\psi_1 = 0.978 > 0.9$, which together imply that $0.9 < \inf\{\psi_1 : 0 < \alpha < 0.25\}$ as required.

Let $0 < \alpha < 0.25$ and $0 < \psi < 0.9$. We now consider a single parameter, θ , and a corresponding estimator $Y \sim N(\theta, 1)$, and show that the probability that a sign-determining non-covering confidence interval is constructed for θ , is no less than $\alpha/2$ for all θ . Formally, let NCI be the event that $CI := \mathcal{C}_{MQC}(Y; \alpha)$ (i) determines the sign, i.e., does not include values of opposite signs and (ii) does not include the true value θ . Then we show that $\Pr_\theta(NCI) \geq \alpha/2$ for all θ . Since for the MQC interval, sign determination occurs if and only if $|Y| \geq \bar{c}$, we have

$$(A.5) \quad \Pr_\theta(NCI) = \Pr_\theta(|Y| \geq \bar{c}, \theta \notin CI).$$

If the confidence interval were obtained simply by inverting the $1 - \alpha$ acceptance regions $A(\theta)$ in (4.1) the event $\theta \notin CI$ could be replaced by $Y \notin A(\theta)$; however, the confidence interval is obtained by taking the convex hull of the inverse set, in which case it is possible that $Y \notin A(\theta)$ and yet $\theta \in CI$. We can overcome this difficulty by considering the “effective” acceptance regions, $\bar{A}(\theta)$, which take into account the fact that the convex hull of $\{\theta : Y \in A(\theta)\}$ is taken, in that $CI = \{\theta : Y \in \bar{A}(\theta)\}$ (here without the convex hull). Denoting by $l(\theta)$ and $u(\theta)$ the lower and upper endpoints of $A(\theta)$, respectively, and denoting by $\bar{l}(\theta)$ and $\bar{u}(\theta)$ the lower and upper ends of $\bar{A}(\theta)$, respectively, it holds that $\bar{l}(\theta) = \max\{u(\tilde{\theta}) : \tilde{\theta} \leq \theta\}$ and $\bar{u}(\theta) = \min\{l(\tilde{\theta}) : \tilde{\theta} \geq \theta\}$.

Explicitly,

$$(A.6) \quad \bar{A}(\theta) = \begin{cases} (-c_{\alpha/2}, c_{\alpha/2}), & \theta = 0 \\ (-\bar{c}, \tilde{c}), & 0 < \theta \leq \tilde{c} - \bar{c} \\ (-\bar{c}, g(\theta)), & \tilde{c} - \bar{c} < \theta \leq \bar{c} + c_{\alpha/2} \\ (\theta - c_{\alpha/2}, \theta + c_{\alpha/2}), & \bar{c} + c_{\alpha/2} < \theta \end{cases}$$

with $\bar{A}(\theta) = -\bar{A}(-\theta)$ for $\theta < 0$ and where $g(\theta) = \theta + F^{-1}\{2 - \alpha - F(\bar{c} + \theta)\}$.

Now we can write

$$(A.7) \quad \Pr_{\theta}(NCI) = \Pr_{\theta}(|Y| \geq \bar{c}, Y \notin \bar{A}(\theta)),$$

and we note that for $0 < \theta < \bar{c} + c_{\alpha/2}$, $(-c, c) \subset \bar{A}(\theta)$, hence $\Pr_{\theta}(NCI) = \Pr_{\theta}(Y \notin \bar{A}(\theta))$. For $\theta = 0$, this is exactly α .

For $0 < \theta < \tilde{c} - \bar{c}$, $\Pr_{\theta}(Y \notin \bar{A}(\theta)) = \Pr_{\theta}(Y \notin (-\bar{c}, \tilde{c}))$, which is minimized at $\theta = (\tilde{c} - \bar{c})/2$. In order that $\Pr_{(\tilde{c}-\bar{c})/2}(Y \notin \bar{A}(\theta))$ be less than $\alpha/2$, in which case $\Pr_{(\tilde{c}-\bar{c})/2}(NCI) < \alpha/2$, it must hold that $\tilde{c} + \bar{c} > 2c_{\alpha/4}$. We claim that this cannot be the case. Hence, for any α , let ψ^* be the value of ψ for which $\tilde{c} + \bar{c} = 2c_{\alpha/4}$. Then for a fixed α , $\psi < \psi^*$ implies that $\tilde{c} + \bar{c} < 2c_{\alpha/4}$. Now, it can be verified that $\lim_{\alpha \rightarrow 0} \psi^* > 0.9$ (but $\lim_{\alpha \rightarrow 0} \psi^* < 0.94$) and that ψ^* is an increasing function of α on $0 < \alpha < 0.25$, which imply that $\psi^* > 0.9$ for all $0 < \alpha < 0.25$. It follows that $\tilde{c} + \bar{c} < 2c_{\alpha/4}$ for all $0 < \alpha < 0.25$, and we conclude that $\Pr(NCI) \geq \alpha/2$ also for $0 < \theta < \tilde{c} - \bar{c}$.

For $\tilde{c} - \bar{c} < \theta \leq \bar{c} + c_{\alpha/2}$, $A(\theta) = \bar{A}(\theta)$, and since $\Pr_{\theta}(Y \in A(\theta)) = 1 - \alpha$, we have that $\Pr_{\theta}(NCI) = \alpha$.

Finally, for $\theta > \bar{c} + c_{\alpha/2}$ we have $\Pr_{\theta}(NCI) = \Pr_{\theta}(|Y| > \bar{c}, |Y| > \theta + c_{\alpha/2}) \geq \alpha/2$.

In any case, $\Pr_{\theta}(NCI)$ does not drop below $\alpha/2$.

To evaluate the FCR, we follow a computation similar to that in BY. Let $0 < q < 0.25$ and $0 < \psi < 0.9$. Denote by $CI_i(\alpha) = \mathcal{C}_{MQC}(Y_i; \alpha)$ a level $1 - \alpha$ MQC interval using parameter ψ , and by $\bar{c}(\alpha) = \Phi^{-1}(1 - \psi \cdot \alpha)$ the value of the quantity \bar{c} associated with it. Furthermore, let $C_k^{(i)} = \{Y^{(i)} : R_{\min}(Y^{(i)}) = k\}$. For the selective-SDCI procedure of Definition 3 $R_{\min} = R_{CI}$, in which case BY show that

$$(A.8) \quad \text{FCR} = \sum_{i=1}^m \sum_{k=1}^m \frac{1}{k} \Pr \left\{ C_k^{(i)}, i \in \mathcal{S}(\mathbf{Y}), \theta_i \notin CI_i \left(\frac{k \cdot q}{m} \right) \right\}.$$

Using the fact that $i \in \mathcal{S}(\mathbf{Y})$ if and only if $|Y_i| \geq \bar{c} \left(\frac{R_{CI} \cdot q}{m} \right)$, we can replace

the right hand side of the last equality by

$$(A.9) = \sum_{i=1}^m \sum_{k=1}^m \frac{1}{k} \Pr \left\{ C_k^{(i)}, |Y_i| \geq \bar{c} \left(\frac{k \cdot q}{m} \right), \theta_i \notin CI_i \left(\frac{k \cdot q}{m} \right) \right\}$$

$$(A.10) = \sum_{i=1}^m \sum_{k=1}^m \frac{1}{k} \Pr \left\{ C_k^{(i)} \right\} \times \Pr \left\{ |Y_i| \geq \bar{c} \left(\frac{k \cdot q}{m} \right), \theta_i \notin CI_i \left(\frac{k \cdot q}{m} \right) \right\}$$

$$(A.11) \geq \sum_{i=1}^m \sum_{k=1}^m \frac{1}{k} \Pr \left\{ C_k^{(i)} \right\} \times \frac{kq}{2m}$$

$$(A.12) = \frac{q}{2}$$

where inequality (A.11) follows from the preceding part of the proof as $\frac{k \cdot q}{m} \leq q < 0.25$.

A.9. Proof of Theorem 3. Beginning with an expression for FCR as appears in Benjamini and Yekutieli [6],

$$\begin{aligned} \text{FCR}(\widetilde{CI}_{\bullet}; \mathcal{S}; q_1, q_2) &= \sum_{r=1}^m \sum_{i=1}^m \frac{1}{r} \cdot \Pr(|\mathcal{S}_2| = r, N\widetilde{CI}_i) \\ &= \sum_{r=1}^m \sum_{i=1}^m \frac{1}{r} \cdot \{ \Pr(|\mathcal{S}_2| = r, N\widetilde{CI}_{i1}) + \Pr(|\mathcal{S}_2| = r, \theta_{i1} \in CI_{i1}, N\widetilde{CI}_{i2}) \} \\ &= \sum_{r=1}^m \sum_{i=1}^m \frac{1}{r} \cdot \Pr(|\mathcal{S}_2| = r, i \in \mathcal{S}_2, \theta_{i1} \notin CI_{i1}) \\ &\quad + \sum_{r=1}^m \sum_{i=1}^m \frac{1}{r} \cdot \Pr(|\mathcal{S}_2| = r, \theta_{i1} \in CI_{i1}, N\widetilde{CI}_{i2}) \\ &= \sum_{r=1}^m \sum_{i=1}^m \frac{1}{r} \cdot \Pr(\theta_{i1} \notin CI_{i1}) \cdot \Pr(|\mathcal{S}_2| = r, i \in \mathcal{S}_2) \\ &\quad + \Pr(\theta_{i1} \in CI_{i1}) \cdot \sum_{r=1}^m \sum_{i=1}^m \frac{1}{r} \cdot \Pr(|\mathcal{S}_2| = r, N\widetilde{CI}_{i2}) \\ &= q_1 \cdot \sum_{r=1}^m \frac{1}{r} \cdot \sum_{i=1}^m \Pr(|\mathcal{S}_2| = r, i \in \mathcal{S}_2) + (1 - q_1) \cdot \sum_{r=1}^m \sum_{i=1}^m \frac{1}{r} \cdot \Pr(|\mathcal{S}_2| = r, N\widetilde{CI}_{i2}) \end{aligned}$$

To complete the proof, note that for any \mathcal{S}_2 , $\sum_{i=1}^m \Pr(|\mathcal{S}_2| = r, i \in \mathcal{S}_2) = r \cdot \Pr(|\mathcal{S}_2| = r)$, and that $\text{FCR}(CI_{\bullet 2}; \mathcal{S}_2; q_2) = \sum_{r=1}^m \sum_{i=1}^m \Pr(|\mathcal{S}_2| = r, N\widetilde{CI}_{i2})/r$.

ACKNOWLEDGEMENTS

Section 7 makes use of data generated by the Wellcome Trust Case Control Consortium. A full list of the investigators who contributed to the generation of the data is available from www.wtccc.org.uk. Funding for the project was provided by the Wellcome Trust under award 076113.

REFERENCES

- [1] Alan Agresti. Categorical data analysis. *A John Wiley and Sons, Inc. Publication, Hoboken, New Jersey, USA*, 2002.
- [2] Rina Foygel Barber and Emmanuel J Candes. A knockoff filter for high-dimensional selective inference. *arXiv preprint arXiv:1602.03574*, 2016.
- [3] Y Benjamini, Y Hochberg, and Y Kling. False discovery rate control in pairwise comparisons. *Research Paper, Department of Statistics and Operations Research, Tel Aviv University*, 1993.
- [4] Y Benjamini, Y Hochberg, and PB Stark. Confidence intervals with more power to determine the sign: Two ends constrain the means. *Journal of the American Statistical Association*, 93(441):309–317, 1998.
- [5] Yoav Benjamini and Daniel Yekutieli. The control of the false discovery rate in multiple testing under dependency. *Annals of statistics*, pages 1165–1188, 2001.
- [6] Yoav Benjamini and Daniel Yekutieli. False discovery rate-adjusted multiple confidence intervals for selected parameters. *Journal of the American Statistical Association*, 100(469):71–81, 2005.
- [7] Robert Bohrer. Multiple three-decision rules for parametric signs. *Journal of the American Statistical Association*, 74(366a):432–437, 1979.
- [8] Robert Bohrer and Mark J Schervish. An optimal multiple decision rule for signs of parameters. *Proceedings of the National Academy of Sciences*, 77(1):52–56, 1980.
- [9] Paul R Burton, David G Clayton, Lon R Cardon, Nick Craddock, Panos Deloukas, Audrey Duncanson, Dominic P Kwiatkowski, Mark I McCarthy, Willem H Ouwehand, Nilesh J Samani, et al. Genome-wide association study of 14,000 cases of seven common diseases and 3,000 shared controls. *Nature*, 447(7145):661–678, 2007.
- [10] Geraldine M Clarke, Carl A Anderson, Fredrik H Pettersson, Lon R Cardon, Andrew P Morris, and Krina T Zondervan. Basic statistical analysis in genetic case-control studies. *Nature protocols*, 6(2):121–133, 2011.
- [11] H Finner. Two-sided tests and one-sided confidence bounds. *The Annals of Statistics*, pages 1502–1516, 1994.
- [12] William Fithian, Dennis Sun, and Jonathan Taylor. Optimal inference after model selection. *arXiv preprint arXiv:1410.2597*, 2014.
- [13] JL Hodges and EL Lehmann. Testing the approximate validity of statistical hypotheses. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 261–268, 1954.
- [14] Jason D Lee, Dennis L Sun, Yuekai Sun, and Jonathan E Taylor. Exact post-selection inference with the lasso. *arXiv preprint arXiv:1311.6238*, 2013.
- [15] Jerzy Neyman, SS Gupta, and DS Moore. Synergistic effects and the corresponding optimal version of the multiple comparison problem. *Statistical Decision Theory and Related Topics*, 2:297–311, 1977.
- [16] Thomas Nichols and Satoru Hayasaka. Controlling the familywise error rate in functional neuroimaging: a comparative review. *Statistical methods in medical research*, 12(5):419–446, 2003.

- [17] Russell A Poldrack and Jeanette A Mumford. Independence in roi analysis: where is the voodoo? *Social Cognitive and Affective Neuroscience*, 4(2):208–213, 2009.
- [18] John W Pratt. Length of confidence intervals. *Journal of the American Statistical Association*, 56(295):549–567, 1961.
- [19] Jonathan D Rosenblatt and Yoav Benjamini. Selective correlations - not voodoo. *To appear, Neuroimage*, 2014.
- [20] Yuekai Sun and Jonathan E Taylor. Valid post-correction inference for censored regression problems. *arXiv preprint arXiv:1403.3457*, 2014.
- [21] Jonathan Taylor, Richard Lockhart, Ryan J Tibshirani, and Robert Tibshirani. Exact post-selection inference for forward stepwise and least angle regression. *arXiv preprint arXiv:1401.3889*, 7, 2014.
- [22] Xiaoying Tian and Jonathan E Taylor. Selective inference with a randomized response. *arXiv preprint arXiv:1507.06739*, 2015.
- [23] Sabrina M Tom, Craig R Fox, Christopher Trepel, and Russell A Poldrack. The neural basis of loss aversion in decision-making under risk. *Science*, 315(5811):515–518, 2007.
- [24] Asaf Weinstein, William Fithian, and Yoav Benjamini. Selection adjusted confidence intervals with more power to determine the sign. *Journal of the American Statistical Association*, 108(501):165–176, 2013.
- [25] Hua Zhong and Ross L Prentice. Bias-reduced estimators and confidence intervals for odds ratios in genome-wide association studies. *Biostatistics*, 9(4):621–634, 2008.

DEPARTMENT OF STATISTICS
 SEQUOIA HALL
 390 SERRA MALL, STANFORD UNIVERSITY
 STANFORD, CA 94305-4065
 E-MAIL: asafw@stanford.edu

DEPARTMENT OF STATISTICS AND
 OPERATIONS RESEARCH
 SCHOOL OF MATHEMATICAL SCIENCES
 TEL AVIV UNIVERSITY
 RAMAT AVIV, TEL AVIV 69978, ISRAEL
 E-MAIL: yekutiel@post.tau.ac.il